

(19) World Intellectual Property Organization
International Bureau(43) International Publication Date
27 December 2002 (27.12.2002)

PCT

(10) International Publication Number
WO 02/103577 A2

(51) International Patent Classification⁷: G06F 17/30 (74) Agent: STROZIER, Robert, W.; 2925 Briarpark Drive, Suite 930, Houston, TX 77042 (US).

(21) International Application Number: PCT/US02/19541 (81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZM, ZW.

(22) International Filing Date: 19 June 2002 (19.06.2002) (84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

(25) Filing Language: English (26) Publication Language: English

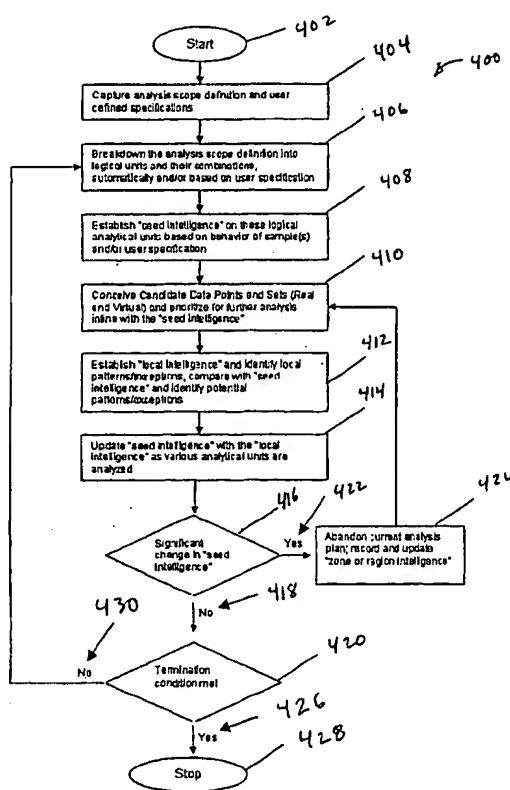
(30) Priority Data: 60/299,243 19 June 2001 (19.06.2001) US

(71) Applicant (*for all designated States except US*): POLYVISTA, INC [US/US]; 1222 Ridgeley Drive, Houston, TX 77055 (US).

(72) Inventors; and (75) Inventors/Applicants (*for US only*): ANWAR, Mohammed, Shahbaz [US/US]; 1222 Ridgeley Drive, Houston, TX 77055 (US). DAHALE, Venkatesh [IN/US]; 5801 Spring Valley #1008, Dallas, TX 75240 (US).

[Continued on next page]

(54) Title: A METHOD FOR EXCEPTIONS DETECTION IN N-DIMENSIONAL DATA SETS WITH FAST CONVERGENCE



(57) **Abstract:** A methodology is described for enhancing data mining processing using virtual database hierarchical constructs, that have dimensionality structure designed for improved data handling by data mining routine or algorithms. The methodology also includes static and/or dynamic data binning routines. The binning routines coupled with the virtual hierarchical constructs provide improved data anomaly detection and enhanced user directed query and data analysis functionality.



Published:

- *without international search report and to be republished upon receipt of that report*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

TITLE: A METHOD FOR EXCEPTIONS DETECTION IN N-DIMENSIONAL DATA SETS WITH FAST CONVERGENCE

INVENTOR: Mohammed S. Anwar and Venkatesh Dahale

ASSIGNEE: PolyVista, Inc., a Corporation of the State of Delaware, United States

RELATED APPLICATIONS

This application claims provisional priority to United States Provisional Patent Application Serial No. 60/299,243 filed 19 June 2001.

BACKGROUND OF THE INVENTION

5

1. Field of Invention

[0001] The present invention relates to a method for detecting exceptions or meta exceptions and/or identifying patterns in base aggregated data and/or automatically and virtually generating data points based on base aggregated data to facilitate rapid data mining in large n-dimensional datasets, especially datasets associated with OLAP cubes.

10

[0002] More particularly, the present invention relates to method for detecting data anomalies, exceptions or meta exceptions and/or identifying patterns in base aggregated data and/or automatically and virtually generating data points based on base aggregated data including the steps of selecting at least one multi-dimensional dataset and at least one measure associated with the data variables; constructing a virtual (imaginary) database schema from a native database schema of the dataset to reduce the dimensionality of the data, while maintaining the measure or producing a Meta measure of more than one measure; selecting a limited number of data variables from the native schema of the dataset; creating an initial global rule describing the behavior of the measure with respect to the limited number of data variables; determining regions of data that would violate the initial global rule; selecting one of the regions; searching the dataset for data that falls within the selected region to form an exception dataset; and reporting the exception dataset.

15

2. Description of the Related Art

[0003] Computer technology continues to allow the storage of ever increasing amount of data. This data storage explosion has given rise to new database technology for storing and retrieving data. However, the interfaces between the user and algorithms designed to delve into the data to find hidden relationships, anomalies, trends or the like is lagging far behind.

20

[0004] Large databases are usually aggregated and summarized along pre-defined, frequently

25

used taxonomy or hierarchy in Multi-Dimensional OLAP databases or in Relational databases to improve ad-hoc query performance in accessing aggregated information.

[0005] While such structures facilitates ad-hoc querying and cross-tab reporting, which remains the predominant objective for such systems, the search for exceptions and patterns usually remains limited along and across these pre-defined taxonomies and based on biased numeric thresholds of one or more measures considered individually or in limited combinations.

[0006] Many data mining techniques exist, which may be used to identify patterns or detect exceptions by processing data at a pre-defined level of granularity and across measure combination(s) (Analysis Context). However, such automated techniques suffer from the same drawbacks as the manual techniques, in that they lack intelligence to self determine Analysis Context and progress across and/or switch between Analysis Contexts, given an initial analysis scope (Search Path).

[0007] One way of improving pattern identification and exception detection is to device a way to overcome "Hard Boundaries" imposed by pre-aggregated structures, while utilizing the aggregated values where possible and by making the exception and pattern detection methods intelligent enough to self specify Analysis Context and Search Path, while providing provisions to integrate human intelligence.

[0008] Thus, there is a need in the art for a methodology for increasing data mining of the data stored in databases, especially in the area of anomaly identification.

SUMMARY OF THE INVENTION

[0009] The present invention provides a method implemented on a computer to alter a dimensionality of a multi-dimensional database hierarchical structure, iteratively and dynamically, to enhance, increase and/or make more efficient data processing so that qualified data points are made available to various data mining algorithms. The virtual (imaginary) alteration of the dimensionality of the database structure can be to increase or decrease the dimensionality of any part of the database structure of interest, *i.e.*, to alter the Dimensions, the Levels, the Members, and/or the Measures depending on the requirements of the data mining algorithm, the user query and/or the form of the information sought from the query. The qualified data points can be Crosstabs, Crossjoins, Meta Exceptions or the like. The data mining routine can be any routine that is designed to detect patterns and/or exceptions in multi-dimensional space; with an ability to specify/modify the dimensionality

and/or data point qualifications – inline with algorithm/process/ user requirements. In addition, the special constructs can also be utilized for user defined ad-hoc reporting.

[0010] The present invention provides a computer having stored thereon code of the methodology described above.

5 [0011] The present invention provides a computer readable medium having stored thereon code of the methodology described above.

[0012] The present invention also provides an analysis wizard records the initial definition of search universe and any user defined customization to the pre-aggregated structures, thus accommodating any existing aggregated structures without imposing any special structural requirements.

10 [0013] The method of this invention provides for initial routines that study the major behavior of a sample dataset allowing generation of a "seed intelligence" or global intelligence or rule describing the major behavior of the measure(s). The seed intelligence is used to create new virtual data points and/or crosstabs and to create Analysis Context along and across previously existing data points and virtual data points, which represent data regions which represent data values that would violate the seed intelligence and prioritizing of such candidate Analysis Contexts towards converging on anomaly or exception quickly. As the analysis progresses through the various candidate Analysis Contexts, the "Seed Intelligence" is constantly revised and the candidate Analysis Contexts are re-prioritized and/or revised.

15 The virtual data point members are created by the algorithm based on both - the measure value thresholds, member ranges and qualified member lists in existing taxonomies. This helps in both - improving scalability and in fine tuning the convergence to anomaly or patterns quickly. Local behavior as well as revising overall behavior is used for detecting anomalies and pattern, exceptions and patterns are made available for user perusal as they are detected. Exceptions and patterns are presented within the prevailing Analysis Context in easy to understand form.

20 [0014] The present invention, thus, provides a method for detecting exceptions or meta exceptions and/or identifying patterns in base aggregated data and/or automatically and virtually generating data points based on base aggregated data to facilitate rapid data mining in large n-dimensional datasets, especially datasets associated with OLAP cubes.

25 [0015] The present invention also provides a method for detecting data anomalies, exceptions or meta exceptions and/or identifying patterns in base aggregated data and/or automatically

and virtually generating data points based on base aggregated data including the steps of selecting at least one multi-dimensional dataset, preferably in the form of an OLAP cube, and at least one measure associated with the data dimension in the dataset; capturing a scope of analysis and constraints from a user; constructing a virtual database schema from a native database schema of the dataset to reduce or expand the dimensionality of the dataset as a whole or in regions of interest, while maintaining the associated measure or producing a composite measure from more than one measure; selecting a limited number of data values from the entire dataset or the part of interest; creating an initial global rule, "seed intelligence" describing the behavior of the measure with respect to the selected, limited number of data values; determining data regions that would violate the initial global rule; prioritizing the regions; searching the dataset for data that satisfies the initial seed intelligence and that falls within the regions forming regional datasets; and reporting the regional datasets.

[0016] The present invention also provides a method for detecting data anomalies, exceptions or meta exceptions and/or identifying patterns in base aggregated data and/or automatically and virtually generating data points based on base aggregated data including the steps of selecting at least one multi-dimensional dataset, preferably in the form of an OLAP cube, and at least one measure associated with the data dimension in the dataset; capturing a scope of analysis and constraints from a user; constructing a virtual database schema from a native database schema of the dataset to reduce or expand the dimensionality of the dataset as a whole or in regions of interest, while maintaining the associated measure or producing a composite measure from more than one measure; selecting a limited number of data values from the entire dataset or the part of interest; creating an initial global rule, "seed intelligence" describing the behavior of the measure with respect to the selected, limited number of data values; determining data regions that would violate the initial global rule; prioritizing the regions; searching the dataset for data that satisfies the initial seed intelligence forming a compliance dataset and that falls within the regions forming regional exception datasets; if the regional exception datasets are not null (empty) or do not contain too few data points to support statistical analysis, creating regional intelligence or local intelligence; determining datapoints within each regional exception dataset that represent exceptions to the local intelligence; and reporting the results.

[0017] The present invention also provides a method for detecting data anomalies, exceptions or meta exceptions and/or identifying patterns in base aggregated data and/or automatically

and virtually generating data points based on base aggregated data including the steps of selecting at least one multi-dimensional dataset, preferably in the form of an OLAP cube, and at least one measure associated with the data dimension in the dataset; capturing a scope of analysis and constraints from a user; constructing a virtual database schema from a native database schema of the dataset to reduce or expand the dimensionality of the dataset as a whole or in regions of interest, while maintaining the associated measure or producing a composite measure from more than one measure; selecting a limited number of data values from the entire dataset or the part of interest; creating an initial global rule, "seed intelligence" describing the behavior of the measure with respect to the selected, limited number of data values; determining data regions that would violate the initial global rule; prioritizing the regions; searching the dataset for data that satisfies the initial seed intelligence forming a compliance dataset and that falls within the regions forming regional exception datasets; if the regional exception datasets are not null (empty) or do not contain too few data points to support statistical analysis, creating regional intelligence or local intelligence; determining datapoints within each regional exception dataset that represent exceptions to the local intelligence; update the initial seed intelligence with the local intelligences properly weighted to form an updated seed intelligence; comparing the updated seed intelligence; if the updated seed intelligence is significantly different from the initial seed intelligence, replacing the initial seed intelligence with the updated seed intelligence; repeating the previous three steps, until there is no significant change between the seed intelligence from the previous iteration and this iteration; and reporting the results. The method can also include a final test to determine whether a termination condition has been met, where failure to meet the condition would restart the analysis construction of the scope of analysis step and the method steps would be continued until the condition is met.

[0018] The present invention also provides a method for constructing an intelligence models including an overall or global intelligence and local intelligences using the methods set forth above, which generates the intelligences from the analysis of data in multidimensional databases, relational or OLAP, and in the use the intelligence model to predict further data behavior.

[0019] The present invention also provides a method for constructing libraries of intelligence models, each model including an overall or global intelligence and local intelligences using the methods set forth above, which generates the intelligences from the analysis of data in

multidimensional databases, relational or OLAP.

[0020] The present invention also provides a method for using the library of intelligence models to classify data behavior and as a tool for predicting the behavior of classified data, and in the use the intelligence models to predict further data behavior.

5 [0021] The present invention provides a computer having stored thereon code corresponding to the above provided methods.

[0022] The present invention provides a computer readable medium having stored thereon code corresponding to the above provided methods.

DESCRIPTION OF THE DRAWINGS

10 [0023] The invention can be better understood with reference to the following detailed description together with the appended illustrative drawings in which like elements are numbered the same:

[0024] Figures 1A-D depict the structure of OLAP databases showing dimensions, measures and values and illustrating the formation of composite measure - dimensional reduction and illustrating the identification of exception and meta exception candidate data regions;

15 [0025] Figures 2A-D illustrate a wizard for defining analysis scope and identifying exception candidate data regions;

[0026] Figures 3A-D illustrate the on the fly binning process and the results derived therefrom;

20 [0027] Figure 4 depicts a conceptual flowchart of a preferred method of this invention, which illustrates an iterative method for detecting exceptions and patterns in a specified analysis scope through guided analysis of real and virtual data points and/or crosstabs in an OLAP cube;

25 [0028] Figure 5 shows a cube having defined with a dimensionality of four dimensions and one measure;

[0029] Figure 6 depicts the construction of a composite measure in the cube of Figure 4, the Performance Monitors (Metrics) are arranged in a dimension, the tuple (Member of Performance Monitor Dimension, Member of Measures Dimension) in accord with step 306 of Figure 3, a wizard helps in the specification and customization of any special structures;

30 [0030] Figure 7 shows an example of a wizard, which records user inputs related to a dimension level listing the measure metrics. The user selects the members, which will be studied as composite measures from the Performance Monitors dimension. The user selects

“Memory-Bytes Available” and “Memory-Pages Per Sec” members from the Performance Monitor dimension.

[0031] Figure 8 shows the next screen of the wizard, which allows for the specifying the measures that quantify the value of Performance Monitor members;

5 [0032] Figure 9 depicts a window showing the results of the operation of the wizard of Figure 7;

[0033] Figure 10 depicts the result of dimension reduction by concatenation of measures into the member dimensions;

10 [0034] Figure 11 depicts a screen showing the next step of the wizard operation, where a sample population is polled to allow construction of a seed intelligence or an initial guess of a global rule defining the relationship between the members being correlated;

[0035] Figure 12 depicts a window showing the constructed relationship or seed intelligence from the sample population;

15 [0036] Figure 13 is a plot graphically depicting the seed intelligence as a straight line with negative slope, a grid binning the plot into nine bin valued regions, and the identification of exception regions shaded for the sake of highlighting;

[0037] Figure 14 is a plot graphically depicting a preferred method for determining local intelligence and local intelligence exceptions involving binning data points within the exception regions on a finer scale;

20 [0038] Figure 15 depicts a screen showing crosstab results from MDX code based on the binning process illustrated in Figure 14;

[0039] Figure 16 depicts a screen showing crosstab results from the iterative analysis, where global (“seed intelligence”), regional and local behavior of composite measure is determined and revised.;

25 [0040] Figure 17 depicts a screen showing crosstab results of a hybrid analysis;

[0041] Figure 18 depicts a binning configuration for testing the global intelligence, and identifying and testing local intelligence;

[0042] Figure 19 depicts an iterative processes for analyzing data points within each bin of Figure 18;

30 [0043] Figure 20 depicts a plot of the results of the iterative processes and the identification of an exception within the intelligences;

[0044] Figure 21A and Figure 21B depict a plot showing the seed intelligence and deviations

from the seed intelligence analyzed in a step wise binning process, where first Memory Paging binned slices are analyzed followed by Memory Availability binned slice analysis; and

[0045] Figure 22 depicts a screen of an analysis which has resulted in the confirmation of the seed intelligence.

DETAILED DESCRIPTION OF THE INVENTION

[0046] The inventors have developed a methodology for facilitating the identification of exceptions or anomalies in data via the construction of global rules from a sample selection of data in a multidimensional dataset and for the identification of regions of data that do not obey the rule (exceptions) and the construction of local rules to identify exceptions to the local rules. The inventors have found that this methodology has the following benefits and applications: (1) eliminates the need to create Binned Dimensions with Pre-Defined Intervals-Bins (Stored or Virtual) in OLAP cubes; (2) provides a new way of studying the interaction between Binned Variables in n-Dimensional Space; (3) provides a new way to detect anomalies (Dimensional Context associated with Data Bin Context) based on the concept of Exception and Meta-Exception - Tupled Bins with anomalous associated data; (4) provides a new way to converge on the Meta-Exceptions faster; (5) accommodates special cube structures and measures constructs. These data constructs, when organized in Categorical or Regular Dimensions, can be tupled with Measures to define Meta-Measure or Composite Measure along which the dimensional crosstabs can be reported and/or analyzed for exceptions.

[0047] The present invention relates broadly to a system for finding global and local data patterns and exceptions to both the global pattern and the local pattern, where the system includes an analysis scope capture and definition module, a breakdown module for breaking the analysis scope into logical units or combinations of logical units, a seed intelligence module that determines a seed intelligence (global rule) from a limited data selection from the data to be analyzed; a determine exception candidate region modules where regions of data which would violate the seed intelligence are identified, prioritized and analyzed inline with the analysis of the seed intelligence guess, a determine local intelligence and identify local intelligence exceptions and compare the local intelligence to the seed intelligence, a create an updated seed intelligence module, where the updated seed intelligence and test the updated seed intelligence against the current seed intelligence and repeat the analysis until

the updated seed intelligence and current seed intelligence differ by only an insignificant amount.

[0048] The present invention also broadly relates to a method for finding global and local intelligences quickly including the steps of capturing an analysis scope, a breakdown of the analysis scope into one or more logical units or combinations thereof, optionally specifying constraints on the analysis scope; establishing a seed intelligence from a sample data population, from user input or a combination of data sampling and user input, identifying data regions that represent exceptions to the seed intelligence, establishing local intelligence in each exception region, if non empty, updating seed intelligence with local intelligences or forming a composite intelligence of an updated seed intelligence and local intelligences; testing to determine if the seed intelligence or composite intelligence from the last cycle is significantly different than the seed intelligence or composite intelligence of this cycle, exiting changes are insignificant or returning to the identifying step if significant changes occurred for iteration until convergence is achieved. The method can also include a termination test. After convergence, the method will have constructed a consistent intelligence, seed or composite, for describing the data behavior and will have identified exceptional regions, local intelligence associated with the regions and exceptions to the local rules.

[0049] The present invention also relates broadly to the construction of intelligences, seed, local and/or composite, for the construction of model for predicting data behavior. The intelligences can also be pooled into a library for even faster trend and pattern analysis of n-dimensional datasets.

[0050] The methodology of the present invention is ideally suited for finding data exceptions, global and local intelligences or data patterns and composite intelligences or data patterns - mixtures of global and local data patterns or intelligences in data with many dimensions contained in any type of database, but preferentially contained within an OLAP database. The methods of this invention are ideally suited for the analysis of any type of multidimensional data including, without limitation, operational data, manufacturing data, financial data, currency exchange data, human behavioral data, medical data, regulatory data, legal data, or any other data that have many dimensions (members) and at least one measure (value).

[0051] The methods of this invention allow dimensional manipulations of the original database schema without having to change the original database schema. Thus, the method

-10-

creation of expanded or reduced database schema to construct database schemas that do not correspond to the physical database schema of the databases being analyzed. These computer constructed database schema can be generated by the method, specified by the user or any combination of method generated and user specified constructions. The imagined or constructed database are designed to improve the efficiency and speed in generating global and local data patterns - intelligences, through an iterative or recursive method to refine the intelligences until they are self-consistent – do not vary significantly from one cycle to the next. The meaning of significant, of course, may change from analysis to analysis and may even be user defined. However, the term generally means that intelligences from two consecutive cycles differ by less than about 20% at each point along the graphical or mathematical representation of the intelligence, global, local and/or composite, preferably, the difference is less than about 15% at each point, particularly, less than about 10% at each point, and more particularly, less than about 5% at each point, with an ultimate goal being less than 1% at each point. Of course, the smaller the acceptable difference, the longer the process will take to converge. Thus, for gross analysis, a 20% or greater difference may be acceptable; while for a detailed analysis, 1% or less may be acceptable.

[0052] This invention also relates to a method for automatically, interactively and dynamically generating candidate virtual data points and selecting real data points (Crosstabs, Crossjoins, Meta Measures and Meta Exceptions) in OLAP and RDBMS databases for a specified analysis scope; *a priori* and *post-priori* application of statistical, data mining techniques to prune and/or prioritize candidate data points per the analysis objective; application of statistical and data mining techniques to identify exceptions in or patterns across candidate data points and potentially, interactively revise candidate data point generation definition and prioritization to fine tune and expedite exception or pattern detection.

[0053] The method can also include generating candidate virtual data points includes virtual alteration of dimensionality (Splitting or Merging Dimensions, Levels, Members, Measures in the multi-dimensional database, Columns and Rows of Tables in Relational Databases. The method can also include generating candidate virtual data points is based on analysis definition and user defined, algorithm defined –static and/or dynamic thresholds and context.

30 The method can also include generating candidate virtual and real data points are prioritized for and interactively subjected to statistical and data mining routines to detect exceptions and identify patterns. The method can also include detecting exceptions and patterns are

presented in the context of combination of real and virtual data points, which includes prevailing thresholds and conditions. The method can also include combining real and virtual data points are utilized for user defined ad-hoc reporting or analysis.

[0054] This invention also relates to a method to alter dimensionality (increase/decrease) (Dimensions, Levels, Members, Measures) of the multi-dimensional database, interactively and dynamically, such that qualified data points (Crosstabs, Crossjoins, Meta Exceptions) are made available to various algorithms (that detect patterns and/or exceptions in multi-dimensional space; with an ability to specify/modify the dimensionality and/or data point qualifications – inline with algorithm/process/ user requirements) for efficient processing.

5 In addition, the special constructs can also be utilized for user defined ad-hoc reporting. An application of the above method for crosstab qualifications (based on user defined, algorithm defined - static or dynamic thresholds).

[0055] This invention also provide a method that utilizes the above two concepts to converge on multi-variant anomaly fast – Sizing the problem based on Meta-Exceptions, while simultaneously selecting viable/optimal candidates (Candidates can be one or more dimensions (and there combinations) and/or one or more members (and there aggregated combinations).

[0056] The methods of this invention present a unique way of displaying the detected results such that anomalies are presented in a Cause and Effect relationship that include various prevailing thresholds and conditions used by the algorithm.

Introduction

[0057] As an introduction to the methodology of this invention, in an OLAP cube, data lies at the intersection of dimensional members. Looking at Figure 1A, an OLAP cube having four dimensions is shown schematically in a display window 100. The display window 100 includes columns 102a-d for each of the four dimension: Dim A, Dim B, Dim C, and Dim D, a column 104 for the Dim Measure and a column 106 for the intersection value. The display window 100 also includes a header row 108 with headers boxes 110 and ten data rows 112.

[0058] In an OLAP environment, data values are associated with intersections of dimension members and measures. Ad-Hoc reporting entails viewing data values that lie at the intersections of desired dimensional members. Ad-Hoc analysis entails viewing dimensional members that intersect to yield desired data value. Using Ad-Hoc analysis, the search for

-12-

unusual data or exceptional data values with respect to expected or predicted trend in the data values is described below.

Important Facts

[0059] An abnormal intersection value constitutes an exception to a general rule of expected or predicted data behavior. In the methodology of this invention, such exceptions can be applied as filters and can be defined by a composite of dimensional member constructs – a tuple of regular dimensional members and measures and conditional “Intersection Value.” Looking at Figure 1B, a window 120 is shown that includes **Candidates for Exception**; while in Figure 1C, a window 122 is shown that includes **Candidates for Meta Exception**, where the difference between the window 120 and the window 122 lies in the reduction of the dimensionality by merging dimensions – Dim C + Dim D + Dim Measure → Dim CDMeasure..

Sample MDX – Illustrating Meta Exception

[0060] For simplicity, consider the intersection defined by five dimensions as set forth in the following MDX (Multi-Dimensional Expression – a query language for OLAP databases) code:

```

WITH SET [OLAPINTERSECTION] AS
  '{NonEmptyCrossjoin(
    {[Customers].[Country].MEMBERS},
    {[Education Level].[Education Level].MEMBERS},
    {[Gender].[Gender].MEMBERS},
    {[Marital Status].[Marital Status].MEMBERS}
  )}'

  MEMBER [MEASURES].[METAEXCEPTION1] AS
    '([P r o d u c t] . [A l l Products].[Food].[Dairy].[Measures].[Unit Sales])'
  MEMBER [MEASURES].[METAEXCEPTION2] AS
    '([Promotion Media].[All Media].[Bulk Mail].[Measures].[Sales Count])'
  SELECT
    FILTER([OLAPINTERSECTION], /*Intersections where Filter Conditions (1,2) are met*/
      (SUM(([OLAPINTERSECTION].CURRENTMEMBER), [MEASURES].[METAEXCEPTION1])>80)
    OR
      (SUM(
        ([OLAPINTERSECTION].CURRENTMEMBER), [MEASURES].[METAEXCEPTION2]
      )< 8)
    )
  ON COLUMNS
  FROM SALES

```

[0061] Looking at Figure 1D, a display window 130 is shown including a sample meta-

-13-

exception result window, where Dim A is customers, Dim B is education, Dim C is gender, and Dim D is marital status, generated by the above MDX code is shown, where only non empty intersections corresponding to data values where the two Meta Exception rules are met.

[0062] Figures 2A-D illustrate the steps the methodology of this invention uses to specify a search for exceptional data. Looking at Figure 2A, a window is shown for selecting a crosstab measure type. Looking at Figure 2B, a window is shown for selection of transposed measure dimension "Performance Monitor." Looking at Figure 2C, a window is shown with the third step of the exception definition process, where the elements of the transposed measures dimension are displayed and selected. At this step, user constraints can be added between the measures and/or between the transposed measures dimension members. Looking at Figure 2D, a window is shown including the constructed measure and transposed measures dimension and child and descendent measure values in the crosstab columns next to the list of dimensions.

[0063] The MDX code generated and executed using the definition step described above is shown below:

```

WITH SET [OLAPINTERSECTION] AS
  '{NonEmptyCrossjoin(
    {[Hour].[Hour].MEMBERS},
    {[Configuration].[Configuration Name].MEMBERS},
    {[Computer].[Name].MEMBERS}
  )}'
  MEMBER [MEASURES].[METAEXCEPTION1] AS'([Performance Monitor].[All
  Performance Monitor].[Memory-Bytes Available], [Measures].[Sample
  Avg])'
  MEMBER [MEASURES].[METAEXCEPTION2] AS
  '([Performance Monitor].[All Performance Monitor].[Memory-Pages
  per Sec], [Measures].[Sample Avg])'

  SELECT
  FILTER([OLAPINTERSECTION],
    (SUM      ( {[OLAPINTERSECTION].CURRENTMEMBER},
    [MEASURES].[METAEXCEPTION1] )>800000000)
  AND
    (SUM      ( {[OLAPINTERSECTION].CURRENTMEMBER},
    [MEASURES].[METAEXCEPTION2] )< 800000000))
  ON COLUMNS

  FROM PERFORMANCE

```

[0064] The present invention also relies on the concept of data value binning to help simply facilitate exception identification and pattern construction. Binning can reduce the amount of data values to be analyzed and help to augment the local data behavior. The particular type of binning most useful in the application of this invention is so-called "on the fly binning."

On the Fly Binning – Based on Exception and Meta Exceptions Concepts

[0065] Binning can be defined as a process of mapping continuous values into categorical values or bins. A bin is a category, e.g., a series of continuous values 1,2,3,4,5,6,7,8,9,10 can be binned to the following categorical valued bins 1 to 5 and 5 to 10. Binning adds a lot of value in the process of exception detection. Binning can amplify data effects, such that the previously diluted exceptions, which were hard to identify in the entire data population are now easily identified in segments of population (Profiling – Binned Members). Binning can also reduce the effort required for exception detection by providing a sampling approach (Sampling – Binned Sets).

[0066] In the OLAP world, the bins can be defined by the dimensional members and data values. The data values, Val(Tuple), can be binned into absolute or dynamic ranges to form complex bins. For example, the binning can be as simple as:

```
([Measures].[Sales Count] ) > 5000 AND ([Measures].[Sales Count]
) < 5900 )
```

or more complex as:

```
([Customers].[All Customers].[Canada], [Education Level].[All
Education Level].[Bachelors Degree], [Gender].[All Gender].[M],
[Product].[All Products].[Drink].[Alcoholic Beverages].[Beer and
Wine].[Beer], [Measures].[Unit Sales] ) >
2([MEASURES].[METAEXCEPTION1])
AND
([Customers].[All Customers].[Canada], [Education Level].[All
Education Level].[Bachelors Degree], [Gender].[All Gender].[M],
[Product].[All Products].[Drink].[Alcoholic Beverages].[Beer and
Wine].[Beer], [Measures].[Unit Sales] ) <
3([MEASURES].[METAEXCEPTION1])
```

[0067] The bins can be based on equal count per bin, equi-count bins, user defined bins or dynamically set by using outlier identifiers such as standard deviation, average, median, mode, min, max or other statistical functions. As an example of definition and utilization of bins, consider the following MDX code:

```
SELECT
  ORDER(([Product].[Product Name].Members), [Measures].[Unit
  Sales], BASC) ON ROWS,
  {[Measures].[Unit Sales],
  [Measures].[Profit], [Measures].[Sales Count]} ON COLUMNS
  FROM SALES
```

[0068] The observed values for Unit Sales and Profit for the Products range over their represented data values in the dataset.

[0069] To sample/profile the Products based on Unit Sales and Profit to study Sales Count patterns, binning would be performed on the Products based on Unit Sales and Profit Values. For example, consider following MDX statements, which shows a sampling application:

WITH

-15-

```

SET [ProductUnitSales_BIN1] AS 'FILTER({{[Product].[ProductName].Members}, ISNULL([Measures].[Unit Sales])})'
SET [ProductUnitSales_BIN2] AS 'FILTER({{[Product].[ProductName].Members}, (([Measures].[Unit Sales] > 0) AND ([Measures].[Unit Sales] < 100)))}'
SELECT
    UNION({[ProductUnitSales_BIN1]}, {[ProductUnitSales_BIN2]}) ON ROWS, {[Measures].[Unit Sales]}, [Measures].[Profit], [Measures].[Sales Count] ON COLUMNS
FROM SALES

```

[0070] The result of the above operation shows that we can apply the bins in terms of sets, as a bin range may return multiple products in a profile. The operation also shows that multiple bins can be unioned together to form sets of larger bin.

```

WITH
    SET [ProductUnitSales_BIN2] AS 'FILTER({{[Product].[ProductName].Members}, (([Measures].[Unit Sales] > 0) AND ([Measures].[Unit Sales] < 130)))'
    SET [ProductUnitSales_BIN3] AS 'FILTER({{[Product].[ProductName].Members}, (([Measures].[Profit]> 50) AND ([Measures].[Profit]< 80)))'
SELECT
    UNION({[ProductUnitSales_BIN2]}, {[ProductUnitSales_BIN3]}) ON ROWS, {[Measures].[Unit Sales]}, [Measures].[Profit], [Measures].[Sales Count] ON COLUMNS
FROM SALES

```

[0071] The results from the above MDX code presents Products that belong to Bin2 or Bin3 constructed using the UNION function. Alternatively, using the INTERSECT function, the code would present Products that belong to both Bin2 and Bin3

[0072] The method of this invention allows binning without having to key in the MDX code manually. For the data set forth above, the method would bin the Unit Sales and Sales Count at the Relational Level, then build the Bin Dimensions in OLAP cube and finally select the appropriate range dimension in the Dimension Tree. The present invention allows bins to be based on MetaException concepts, which permits binning based exclusively along particular Dimensions and/or Level members.

[0073] One of the application of "On The Fly Binning" as in the above example would be to sample the OLAP data, which can then be subject to further analysis. For example, consider following MDX statements, which show profiling application:

```

WITH
    Member [Product].[ProductUnitSales_BIN1] as
        'AGGREGATE({FILTER({{[Product].[ProductName].Members}, (([Measures].[Profit]> 50) AND ([Measures].[Unit Sales] < 100))})}'
    Member [Product].[ProductUnitSales_BIN2] as
        'AGGREGATE({FILTER({{[Product].[ProductName].Members}, (([Measures].[Profit]> 50) AND ([Measures].[Profit]< 53))})}'
SELECT

```

-16-

```

{ [Product].[ProductUnitSales_BIN1],
[Product].[ProductUnitSales_BIN2])) ON ROWS,
{[Measures].[Unit Sales], [Measures].[Profit],
[Measures].[Sales Count]} ON COLUMNS
FROM SALES

```

[0074] The results of the above MDX codes presents two new members which are aggregates of the specified filter condition. Such members can be used to magnify/concentrate exceptions/patterns that would previously be less visible. Note that each of the bins represents is a Unique Member, so there is no need to use UNION or INTERSECT functions. For example, consider following MDX statements, which show another profiling application:

```

WITH
Member [Product].[ProductUnitSales_BIN1] as
'AGGREGATE({FILTER({{[Product].[Product Name].Members}, (([Measures].[Profit]> 50) AND ([Measures].[Unit Sales]< 100))}))'
Member [Product].[ProductUnitSales_BIN2] as
'AGGREGATE({FILTER({{[Product].[Product Name].Members}, (([Measures].[Profit]> 50) AND ([Measures].[Profit]< 53))}))'
Member [Customers].[CustomersUnitSales_BIN3] as
'AGGREGATE({FILTER({{[Customers].[City].Members}, (([Measures].[Profit]> 50) AND ([Measures].[Profit]< 100))}))'
Member [Customers].[CustomersSalesAverage_BIN4] as
'AGGREGATE({FILTER({{[Customers].[City].Members}, (([Measures].[Sales Average]> 0) AND ([Measures].[Sales Average]< '500)))})'

SELECT
{[Product].[ProductUnitSales_BIN1], [Product].[ProductUnitSales_BIN2]} ON ROWS,
{[Customers].[CustomersUnitSales_BIN3], [CustomersSalesAverage_BIN4]} ON
COLUMNS
FROM SALES
WHERE [Measures].[Unit Sales]

```

[0075] The result of this MDX code is displayed Figure 3A, which shows the values for the 2X2 cross tab with ProductUnitSale_Bin1 and 2 on the rows and CustomersUnitSales-Bin3 and 4 on the columns.

[0076] The above MDX results in a crosstab that shows bins along a particular axis based on separate measure values and that the concept could be utilized towards mere ad-hoc report (which is a less significant application) to complex bin permutations.

[0077] Now, consider some binning examples using a tuple of two or more dimensions:

```

WITH
Member [Product].[ProductStoreTypeUnitSales_BIN1] as
'AGGREGATE({FILTER(NONEMPTYCROSSJOIN({[Product].[Product
Category].Members}, {[Store Type].[Store Type].Members}), (([Measures].[Profit]> 0) AND ([Measures].[Unit Sales]< 100))))}'
Member [Product].[ProductStoreTypeSalesAverage_BIN2] as
'AGGREGATE({FILTER(NONEMPTYCROSSJOIN({[Product].[Product
Category].Members}, {[Store Type].[Store Type].Members}), (([Measures].[Sales Average]> 0) AND ([Measures].[Sales Average]< 75 ))}))'

```

-17-

```

Member [Customers].[CustomersEducationUnitSales_BIN3] as
'AGGREGATE({FILTER(NONEMPTYCROSSJOIN({[Customers].[State
Province].Members}, {[Education Level].[Education
Level].Members}), (( [Measures].[Profit]> 0) AND
([Measures].[Profit]< 100))))'
Member [Customers].[CustomersEducationStoreCost_BIN4] as
'AGGREGATE({FILTER(NONEMPTYCROSSJOIN({[Customers].[State
Province].Members}, {[Education Level].[Education
Level].Members}), (( [Measures].[Store Cost]> 0) AND
([Measures].[Store Cost]< 500))))'
SELECT
{[Product].[ProductStoreTypeUnitSales_BIN1], [Product].[Pr
oductStoreTypeSalesAverage_BIN2]} ON ROWS,
{[Customers].[CustomersEducationUnitSales_BIN3],
[Customers].[CustomersEducationStoreCost_BIN4]} ON COLUMNS
FROM SALES
WHERE [Measures].[Unit Sales]

```

[0078] The result of this MDX code is displayed Figure 3B, which shows the values for the 2X2 cross tab with ProductStoreTypeUnitSale_Bin1 and ProductStoreTypeSaleAverage_Bin2 on the rows and CustomersEducationUnitSales_Bin3 and CustomersEducationStoreCose_Bin4 on the columns and showing boxed crosstab member of interest.

[0079] The above Binning (Profiling) operation can be made available to other algorithms that work in multi-dimensional space, resulting in the detection of the highlighted cell as exceptional. The method can create more advanced bins by using multidimensional tuples and by using dynamic ranges. Binned Members have significant overhead as aggregations are calculated on the fly; however, the flexibility and the analytical enhancement offsets the increased computational overhead.

[0080] Although the Bins may not benefit from existing aggregations, aggregations can be flexibly created. This method is more efficient than a relational environment and more flexible than a pure OLAP. For example, consider the following simple Bins along Customer, Product and Promotion Dimensions as viewed along the Customer Dimension:

```

Member [Customers].[Customers_BIN1] as
'AGGREGATE({FILTER({[Customers].[City].Members}, (([Measures].[Unit
Sales]>0) AND ([Measures].[Unit Sales]< 50))))', SOLVE_ORDER = 2
Member [Customers].[Customers_BIN2] as
'AGGREGATE({FILTER({[Customers].[City].Members}, (([Measures].[Unit
Sales])>= 50) AND ([Measures].[Unit Sales]< 100))), SOLVE_ORDER = 2

```

as view along the Product Dimension:

```

Member [Product].[Product_BIN1] as
'AGGREGATE({FILTER({[Product].[Product Category].Members}, (([Measures].[Unit Sales]> 0) AND ([Measures].[Unit Sales]< 100))), SOLVE_ORDER = 1
Member [Product].[Product_BIN2] as
'AGGREGATE({FILTER({[Product].[Product Category].Members}, (([Measures].[Unit Sales])>= 100) AND ([Measures].[Unit Sales]< 150))), SOLVE_ORDER = 1

```

-18-

and finally as viewed along the Promotion Dimension:

```

Member      [Promotions].[Promotion_BIN1]      as
'AGGREGATE({FILTER({[Promotions].[Promotion Name].Members},
(([Measures].[Unit Sales]> 0) AND ([Measures].[Unit Sales]< 100))))',
SOLVE_ORDER = 3
Member      [Promotions].[Promotion_BIN2]      as
'AGGREGATE({FILTER({[Promotions].[Promotion Name].Members},
(([Measures].[Unit Sales]>= 100) AND ([Measures].[Unit Sales]<
200))))', SOLVE_ORDER =3

```

[0081] Once the bins are formed, the combination of bins can easily be determined that yield high values – just by getting a crossjoined crosstab:

```

WITH
    Member      [Product].[Product_BIN1]      as
'AGGREGATE({FILTER({[Product].[Product Category].Members},
(([Measures].[Unit Sales]> 0) AND ([Measures].[Unit
Sales]< 100)))), SOLVE_ORDER = 1
    Member      [Product].[Product_BIN2]      as
'AGGREGATE({FILTER({[Product].[Product Category].Members},
(([Measures].[Unit Sales]>= 100) AND ([Measures].[Unit Sales]< 150)))), SOLVE_ORDER = 1
    Member      [Customers].[Customers_BIN1]      as
'AGGREGATE({FILTER({[Customers].[City].Members},
(([Measures].[Unit Sales]>0) AND ([Measures].[Unit Sales]< 50)))), SOLVE_ORDER = 2
    Member      [Customers].[Customers_BIN2]      as
'AGGREGATE({FILTER({[Customers].[City].Members},
(([Measures].[Unit Sales]>= 50) AND ([Measures].[Unit Sales]< 100)))), SOLVE_ORDER = 2
    Member      [Promotions].[Promotion_BIN1]      as
'AGGREGATE({FILTER({[Promotions].[Promotion Name].Members},
(([Measures].[Unit Sales]> 0) AND ([Measures].[Unit Sales]< 100)))), SOLVE_ORDER = 3
    Member      [Promotions].[Promotion_BIN2]      as
'AGGREGATE({FILTER({[Promotions].[Promotion Name].Members},
(([Measures].[Unit Sales]>= 100) AND ([Measures].[Unit Sales]<
200)))), SOLVE_ORDER =3
SELECT
    CROSS JOIN {{[Customers].[Customers_BIN1],
    [Customers].[Customers_BIN2]}, {[Product].[Product_BIN1],
    [Product].[Product_BIN2]}} ON COLUMNS,
    {[Promotions].[Promotion_BIN1], [Promotions].[Promotion_BIN2]} On ROWS
    FROM SALES
    WHERE [Measures].[Unit Sales]

```

[0082] The result of this MDX code is displayed Figure 3C, which shows the values for the 2X2 cross tab with Customer_Bin1 and 2 on the outer rows and their levels Product_Bin1 and 2 on the inner rows and Promotion_Bin1 and Promotion_Bin2 on the columns and showing boxed crosstab member of interest.

[0083] Evaluating this in a single crosstab, however, requires a lot of wait time and results in further complications because of Solve Order. While this can be used for Ad-Hoc reporting alone, the method of this invention preferably uses a Permutation based algorithm, which combines Binned Members from various dimensions and evaluates the combinations that yield

high/low values. Besides just High/Low combinations, these bins can be used as Composite Cases for Market Basket Analysis.

Conceptual Flowchart Illustration of a Preferred Method

[0084] Referring now to Figure 4, a conceptualization 400 of a preferred method of this invention is shown to include a start step 402, which transfers control to a capture step 404, where an analysis scope and user specifications are defined. After the analysis is defined, the defined analysis is broken down into logical units or combination of logical units in a breakdown step 406. The breakdown can be performed automatically, can be based on user specifications or can be a combination of automated breakdown and user defined breakdown or constraints thereon.

[0085] After the defined analysis is broken down, a seed intelligence is generated for the logical units based on the behavior of a data sample, or is defined by the user or is generated and then modified by the user in an establish seed intelligence step 408, where seed intelligence represent a guess at a global rule describing the behavior of the data within the sample. Next, the method 400 identifies and/or constructs and prioritizes data regions that would represent exceptions to or violations of the seed intelligence in a conceive candidate regions step 410. Once the regions have been identified, continued analysis includes identifying data that satisfies the seed intelligence and data that falls within the regions proceeds simultaneously or inline.

[0086] Once the exception regions have been constructed and/or identified, data are then collected that satisfy the seed intelligence and the exception regions, if sufficient data is found in an exception regions, then local intelligence is generated and compared to the seed intelligence in an establish local intelligence step 412. The local intelligence step 412 also identifies data exceptions to the generated local intelligences. After local intelligence analysis and construction and seed intelligence analysis, the seed intelligence is updated with respect both to the data consistence with the former seed intelligence and with respect to the local intelligences relative to a strength of the local intelligences in a update seed intelligence step 414 to produce an updated seed intelligence.

[0087] After the updating, the former and the update seed intelligences are compared in a conditional step 416 to determine whether a significant change in the seed intelligence has occurred. If no significant changes to the former seed intelligence has occurred in this cycle, then control is transferred along a NO branch 418 to a termination test step 420; otherwise control is transferred along a YES branch 422 to an abandon current analysis step 424, where

the former seed intelligence is replaced with the updated and control is then transferred back to the conceive candidate regions step 410. The termination step 420 checks the termination condition and if it is met control is transferred along a YES branch 426 to a stop step 428, where the results of the analysis are reported; otherwise control is transferred along a NO branch 430 to breakdown step 406 and analysis is continued until the termination condition is met.

Application of the Method of this Invention to a Specific Problem

[0088] The method is applied to finding data patterns and exceptions to both global and local rules or intelligences, analysis scope, based on any meaningful combination of members of the Performance Monitor dimension with respect to the Monitor Value measure via the construction of composite measures to accomplish a reduction in dataset dimensionality. The analysis can be further restricted by user specified thresholds and intra-measure and inter-measure conditions that constrain or restrict the analysis scope. Although such constraints on the analysis scope can be specified, in the present example, no thresholds or conditions are specified. Although this application of meta-exceptions and binning based on meta-exceptions is applied to a simple situation considering two composite measures, the method can be applied to n-composite measure problem.

[0089] For the purpose of controlled development, the goal of the application of the method is to analyze data associated with the performance monitor dimension and the monitor value measure and determine a global pattern, local patterns and exceptions and/or meta exceptions. Referring now to Figure 5, a cube having a dimensionality of four with one measure is shown represented by a partial data base schema showing the Performance Dimensions expanded to show the Computer dimension with its members Domain and Computer, the Configuration dimension with its member configuration, the Hour dimension with its members Hour and Min (minute) and the Performance Monitor dimension with its member Performance Monitor and the Measures dimension and its member Monitor Value.

[0090] Figure 6 depicts a crosstab of a composite measure through variable concatenation to reduce the dimensionality of the problem to be analyzed, where the problem has been reduced to finding an exceptional relation between Memory-Bytes Available and Memory-Pages per Sec based on the monitor value measure. In this cube, the Performance Monitors (Metrics) are arranged in a dimension, the tuple (Member of Performance Monitor Dimension, Member of Measures Dimension), which corresponds to step 404 of Figure 4, a wizard helps in the specification and customization of any user defined constructions.

[0091] Figure 7 shows an example of a wizard, which records user inputs related to a dimension level listing the measure metrics, where the user selections are formed into composite measures from the Performance Monitors dimension to analysis. In this figure, the user has selected "Memory-Bytes Available" and "Memory-Pages Per Sec" members from the Performance Monitor dimension.

[0092] Figure 8 shows the next screen of the wizard, which allows for the user to specify the measures that quantify the value of Performance Monitor members and to set constraints on the analysis scope. Thus, the composite measures may now be defined in MDX syntax as follows:

```
MEMBER [Performance Monitor].[Memory-Bytes Available_Monitor
Value] AS '([Performance Monitor].[Memory-Bytes
Available],[Measures].[Monitor Value])'
MEMBER [Performance Monitor].[Memory-Pages per Sec_Monitor Value]
AS '([Performance Monitor].[Memory-Pages per
Sec],[Measures].[Monitor Value])'
```

[0093] The resulting composite measures are defined in the hierarchy of the "Performance Monitor" dimension so that the following query can be constructed in MDX code:

```
WITH
MEMBER [Performance Monitor].[Memory-Bytes Available_Monitor
Value] AS '([Performance Monitor].[Memory-Bytes
Available],[Measures].[Monitor Value])'
MEMBER [Performance Monitor].[Memory-Pages per Sec_Monitor Value]
AS '([Performance Monitor].[Memory-Pages per
Sec],[Measures].[Monitor Value])'
SELECT
{[Performance Monitor].[Memory-Bytes Available_Monitor Value],
[Performance Monitor].[Memory-Pages per Sec_Monitor Value]}
ON COLUMNS
FROM Performance
```

[0094] The above MDX coding yields the results shown in Figure 9. In effect, for the defined analysis scope, the formation of the composite measure has reduced the cube dimensionality to form a virtual cube having the composite measures as dimensions as shown in Figure 10. The new composite measure includes members of the Performance Monitor dimension and serve as regular measure for further analysis. It should be noted that the creation of a composite measure is not necessary and, demonstrates that such customizations are easily and naturally allowed by the method. The composite measure concept is discussed more fully above in the **"On the Fly Binning"** section.

[0095] As shown in Figure 8, the user can specify constraints based on threshold values and conditions. Referring to Figure 11, the next wizard screen is shown, where further constraints related to the analysis search scope may be specified by limiting analysis across specified members, levels and dimensions.

-22-

[0096] For the current example, the selections in shown in Figure 11 indicate that every crosstab evaluated includes members from the Hour dimension at the Hour level. Members from other selected dimensions are included in the crosstabs as required by the iterative routines. Thus, the largest Non Empty search crosstab that can result from the above selections for studying the composite measure can be specified in MDX syntax as follows:

```
NONEMPTYCROSSJOIN(
    {[Computer].[Computer].MEMBERS},
    {[Configuration].[Configuration].MEMBERS},
    {[Hour].[Hour].MEMBERS}
)
```

[0097] Having defined the composite measures and the analysis scope, the next step is to determine the "seed intelligence" for beginning the search for exceptions, local patterns and local pattern exceptions. For the current example, the analysis will focus on identifying exceptions. According to Step 408 of Figure 4, seed intelligence is determined in order to identify candidate crosstabs for analysis. The seed intelligence is constructed using the following MDX query:

```
/* 1) Composite Measure Definition */
WITH
MEMBER [Performance Monitor].[Memory-Bytes Available_Monitor Value] AS '([Performance Monitor].[Memory-Bytes Available],[Measures].[Monitor Value])'
MEMBER [Performance Monitor].[Memory-Pages per Sec_Monitor Value] AS '([Performance Monitor].[Memory-Pages per Sec],[Measures].[Monitor Value])'
/* 2) Sample of High and Low valued members from Hour dimension */
SET [Hour Members] AS
'Union({TopPercent({[Hour].[Hour].members},10, [Performance Monitor].[Memory-Bytes Available_Monitor Value]}), {BottomPercent({[Hour].[Hour].members},10, [Performance Monitor].[Memory-Bytes Available_Monitor Value]}))'
/* 3) Sample of High and Low valued members from Computer dimension */
SET [Computer Members] AS
'Union({TopPercent({[Computer].[Computer].members}, 10, [Performance Monitor].[Memory-Bytes Available_Monitor Value]}), {BottomPercent({[Computer].[Computer].members},10, [Performance Monitor].[Memory-Bytes Available_Monitor Value]}))'
/* 4) All members from Configuration dimension (Few Members) */
SET [Configuration Members] AS
'{[Configuration].[Configuration].members}'
/* 5) Sample of High and Low valued tuples from Hour, Computer and Configuration dimensional combinations */
Set [Scope Members] AS
'Union({TopPercent(NonEmptyCrossjoin({[Configuration].[Configuration].members}, {[Hour].[Hour].members}, {[Computer].[Computer].members}), 10, [Performance Monitor].[Memory-Bytes Available_Monitor Value]}), {BottomPercent(NonEmptyCrossjoin({[Configuration].[Configuration].members}, {[Hour].[Hour].members}, {[Computer].[Computer].members}),10, [Performance
```

-23-

```

Monitor].[Memory-Bytes Available_Monitor Value]))'
/* 6) Correlation between the composite measures over Hour
dimension members */
MEMBER [Performance Monitor].[Correlation_Hours] AS
'CORRELATION({[HourMembers]}, [Performance Monitor].[Memory-Bytes
Available_Monitor Value], [Performance Monitor].[Memory-Pages
per Sec_Monitor Value])'
/* 7) Correlation between the composite measures over Computer
dimension members */
MEMBER [Performance Monitor].[Correlation_Computer] AS
'CORRELATION({[ComputerMembers]}, [Performance Monitor].[Memory-
Bytes Available_Monitor Value], [Performance Monitor].[Memory-
Pages per Sec_Monitor Value])'
/* 8) Correlation between the composite measures over
Configuration dimension members */
MEMBER [Performance Monitor].[Correlation_Configuration] AS
'CORRELATION({[ConfigurationMembers]}, [Performance
Monitor].[Memory-Bytes Available_Monitor Value], [Performance
Monitor].[Memory-Pages per Sec_Monitor Value])'
/* 9) Correlation between the composite measures over all
dimensional tuples */
MEMBER [Performance Monitor].[Correlation_Overall] AS
'CORRELATION({[ScopeMembers]}, [Performance Monitor].[Memory-
Bytes Available_Monitor Value], [Performance Monitor].[Memory-
Pages per Sec_Monitor Value])'
/* 10) Query to return all the correlation values */
Select {[Performance Monitor].[Correlation_Hours], [Performance
Monitor].[Correlation_Computer], [Performance
Monitor].[Correlation_Configuration], [Performance
Monitor].[Correlation_Overall]} ON COLUMNS
FROM Performance

```

[0098] The above MDX code yield the correlated results shown in Figure 12. In the MDX code, sections 2, 3, 4, and 5 relate to selecting the top 10% (or N% depending on sampling needs) and the bottom 10 % (or M% depending on sampling needs) of the members along the selected dimensions, based on one of the composite measures. Different techniques and bias (such as Prominent Members, User Defined Members, Random Selection, *etc.*) can be used to make this sampling more useful, while reducing the resources in deriving such preliminary intelligence.

[0099] Sections 6, 7, 8, and 9 relate to calculating correlations between the composite measures over the sampled members along the selected dimensions. It is important to note that Correlation is one of the many statistical techniques that can be employed here to understand the relationship between the composite measures in the analysis context.

[0100] Further, several permutations and combinations of members can be used to calculate behavior as discussed in step 410 of Figure 4. In that step, the method derives a useful "seed intelligence" which helps to guide the method towards generating and prioritizing data points or crosstabs for exception or pattern detection.

[0101] Referring again to Figure 12, the above determination of the seed intelligence indicates

that the dominant relationship between the composite measures is an inverse relationship between the members based on the measure. Although the method generated an inverse relationship between members, the user could have also defined an initial intelligence from his/her own knowledge of computers. Normally, paging occurs when memory availability is low and paging does not occur or occurs minimally when memory availability is high – an inverse relations. The method is designed to detect situations (Dimensional Contexts) in which this general relationship between available memory and paging is not true. The goal of the method is to identify such situations quickly on a user demand basis, a periodic basis or a continuous basis. Thus, exceptions should occur where Memory Availability and Memory Paging are positively or directly correlated.

[0102] Referring now to Figure 13, a plot 500 of the seed intelligence of composite measure values across the entire selected analysis scope is shown, where the seed intelligence is shown a straight line 502 having a negative slope derived from the analysis shown in Figure 12. The plot 500 has been divided into a grid 504 having nine regions 506 corresponding to low, medium and high values for the two correlated members, Memory Availability on the vertical axis and Memory Paging on the horizontal axis. Because the seed intelligence is an inverse relationships, regions of exceptions 508a&b can be easily identified. Region 508a corresponds to data points that have high values for both members simultaneously, a clear violation of the seed intelligence rule; while region 508b corresponds to data points that have low values for both members simultaneously, another type of clear violation of the seed intelligence. Although the seed intelligence generated in the example is a simple straight line, the same type of candidate regions identification formulation can be used for even very complex curve fit correlations, because the regions where exceptions would be found are as well defined after seed intelligence construction as the seed intelligence is itself.

[0103] Now that the two shaded regions 508a and 508b have been identified, which again represent regions where existence of any data points are potentially contrary to the seed intelligence, the method prioritizes the analysis across the regions to converge on the anomalies quickly. Moreover, conducting analysis by regions or by even smaller binned subregions within the main exception regions, helps to break down a large problem into smaller units; thereby improving scalability.

[0104] Referring now to Figure 14, the concept of binned composite members or the construction of virtual data points is shown as discussed in step 412 of Figure 4. In Figure 14,

-25-

the seed intelligence 502, the grid 504, the nine regions 506 and the exception regions 508a&b are shown. Because the grid 504 is a very course binning construct, the use of these course bins 506 is not a preferred method for developing local intelligence within the two exception regions 508a&b – a single bin 506 cover almost the entire exception regions 508a&b. Thus, the preferred method to develop local intelligence uses a finer binning procedure. In Figure 14, the region 508a is analyzed using bins that are much smaller than the grid bins 506, where a horizontal bar 510 shows the binning along the Memory Availability axis and the vertical bar 512 shows the binning along the Memory Paging axis. The method then determines whether any data falls within an intersection 514 of the two bars 510 and 512. If data is found within the intersection 514 and in adjacent regions as the analysis progresses, then a local intelligence can be constructed as represented by a line segment 516. This same process can be applied to the exception region 508b, which may yield a different local intelligence represented by a line segment 518.

[0105] This methodology can be illustrated by the following MDX query:

```

WITH
/* 1) Bin along Memory Paging Monitor Value Composite Measure */
MEMBER [Performance Monitor].[Memory-Pages per Sec_(0 To 5)] AS
'([Performance Monitor].[Memory-Pages per Sec],[Measures].[Monitor Value])'
/* 2) Bin along Memory Availability Monitor Value Composite
Measure -Scaled Down*/
MEMBER [Performance Monitor].[Memory-Bytes Available_(30 To 60)] AS
'([Performance Monitor].[Memory-Bytes Available],[Measures].[Monitor Value])/1048576'
SET      [BININTERSECTION] AS
'{NONEMPTYCROSSJOIN({[Computer].[Computer].MEMBERS}, {[Configuration].[Configuration].MEMBERS}, {[Hour].[Hour].MEMBERS}))}'
SELECT
FILTER( { [BININTERSECTION]}, (( [Performance Monitor].[Memory-
Pages per Sec_(0 To 5)] >=0 AND [Performance Monitor].[Memory-
Pages per Sec_(0 To 5)] <=5) AND ([Performance Monitor].[Memory-
Bytes Available_(30 To 60)] >=30 AND [Performance
Monitor].[Memory-Bytes Available_(30 To 60)] <=60))) ON COLUMNS,
{[Performance Monitor].[Memory-Pages per Sec_(0 To 5)], [Performance Monitor].[Memory-Bytes Available_(30 To 60)]} ON
ROWS
FROM Performance

```

[0106] "The MDX code yield the results shown in Figure 15, which shows the crosstab definition being presented in terms of dimensional combination and prevailing conditions and thresholds of composite members. This definition protocol aids in the understanding of results when exceptions and/or patterns are detected and presented to the end-user. An example of a crosstab definition in case of pattern detection or hybrid analysis could be crosstab composed

of members of affinity group and affinity thresholds such as support, confidence, improvement etc.

[0107] In the query, "[BININTERSECTION]" is defined by crossjoining all dimensional members in the analysis scope. As the analysis progresses, various crosstabs defined by the binned values of composite measures along the analysis scope are evaluated. The dimensions and composite members can be arranged in several permutations, such that more meaningful crosstabs are generated for analysis based on user and algorithm intelligence. While dimensional permutations in violation of expected behavior are primary exceptions, the same analysis iteration may also be used for analysis using other statistical techniques which may only utilize the dimensional context for analysis to detect secondary exceptions or patterns.

[0108] Referring now to Figure 16, an iteration of analysis is shown where global ("seed intelligence"), regional and local behavior of composite measure is determined and revised. Here the size of the crosstab being subjected to analysis is governed by the size of the composite measure(s) bin range. Regardless of the composite measure bin sizes, the aggregated system has the same workload to evaluate the crosstab qualifying the prevailing thresholds and conditions.

[0109] In small to medium aggregated structures, this may not be of a significant overhead; however for larger structures analysis routines can be structured such that the "data slice" qualification overhead is split into smaller units. For example, consider the analysis routines below:

[0110] 1. Pivot on composite measure(s) and split, "data slice", qualification overhead as described conceptually in the following pseudo MDX code:

```
/* Composite Measure(s) iteration */
For composite measures iteration ((CompositeA Bin (n), CompositeB
Bin (m)...), Bin Permutation (u))
{
    /* "data slice" to be qualified against prevailing
    threshold */
    For dimensional iteration (Level (Dim1 (i), Dim2 (j), Dim3
    (k) ...), Permutation (l))
    {
        /* "data slice" to be qualified being split into
        smaller units*/
        For member set iteration (Set (Dim1 (o), Dim2 (p),
        Dim3 (q)...))
        {
            /* The qualified spit unit of "data slice" is
            analyzed */
            Exception/Pattern      detection      iteration
            (Statistical Technique (t))
        }
    }
}
```

-27-

[0111] 2. Pivot on dimensional iteration(s) and do not split, "data slice", qualification overhead as described conceptually in the following pseudo MDX code:

```
/* "data slice" to be qualified against prevailing threshold */
For dimensional iteration (Level Members (Dim1 (i), Dim2 (j),
Dim3 (k) ...), Permutation (l))
{
    /* Composite Measure(s) iteration */
    For composite measures iteration ((CompositeA Bin (n),
    CompositeB Bin (m)...), Bin Permutation (u))
    {
        /* The qualified spit unit of "data slice" is
        analyzed */
        Exception/Pattern detection iteration (Statistical Technique (t))
    }
}
```

[0112] The Iteration process involves all possible permutations of member sets and composite measure bins. The composite measure bins can be substituted by cluster group list or affinity groups list etc. in case of pattern or hybrid analysis inline with example "a" as set forth in the following MDX code:

WITH

```
/* 1) Bin along Memory Paging Monitor Value Composite Measure */
MEMBER [Performance Monitor].[Memory-Pages per Sec_(0 To 5)] AS
'([Performance Monitor].[Memory-Pages per Sec],[Measures].[Monitor Value])'

/* 2) Bin along Memory Availability Monitor Value Composite
Measure -Scaled Down*/
MEMBER [Performance Monitor].[Memory-Bytes Available_(30 To 60)]
AS
'([Performance Monitor].[Memory-Bytes Available],[Measures].[Monitor Value])/1048576'

/* 3) Current iteration Split Analysis Scope*/
SET [BININTERSECTION] AS
'NONEMPTYCROSSJOIN(Subset({[Computer].[Computer].MEMBERS},100,
50), SubSet({[Hour].[Hour].Members}, 0,50),
Subset({[Configuration].[Configuration].Members}, 0, 50))'

/* 4) Query to return the qualified data slice*/
SELECT
FILTER( { [BININTERSECTION]}, (( [Performance Monitor].[Memory-
Pages per Sec_(0 To 5)] >=0 AND [Performance Monitor].[Memory-
Pages per Sec_(0 To 5)] <=5) AND ([Performance Monitor].[Memory-
Bytes Available_(30 To 60)] >=30 AND [Performance
Monitor].[Memory-Bytes Available_(30 To 60)] <=60))) ON COLUMNS,
{[Performance Monitor].[Memory-Pages per Sec_(0 To
5)], [Performance Monitor].[Memory-Bytes Available_(30 To 60)]} ON
ROWS

FROM Performance
```

[0113] The MDX code yields the results shown in the crosstab of Figure 17. Section 3 of the above MDX shows that we are including 50 Members (variable) at a time from each dimension;

this breaks down the overhead associated with qualifying the data slice across multiple iterations for the purposes of scalability and fast convergence. The number of members included from each dimension is influenced by the dimension size and the number of dimensions in the cube. Members may be selected sequentially, randomly, based on ranking, inter and intra dimensional affinities and/or based on non -sparse member and/or tuple list persisted through earlier iterations and analysis runs. The main goal being to converge on exceptional data slices as efficiently as possible.

[0114] In section 4, we use the binned threshold values of composite measures to qualify the data slice; these thresholds may be determined automatically by the algorithm and/or specified by the user, within the continuum of upper and lower values of these composite measures. The members per dimension and bin range thresholds are also affected by the data characteristics. The size of the analysis data slice being evaluated needs to optimal for statistical technique being employed to identify exceptions and patterns and minimize "false positive" exceptions or patterns.

[0115] The evaluation routines illustrated in example "a" and "b" above, allow for effective use of parallel processing. Further, when an anomaly/pattern is detected, separate independent processes may be launched to further qualify the finding by including other dimensional and measure entities, *etc.*

[0116] In the preceding discussion related to the analysis crosstab creation and analysis guidance, it is important to note that the algorithmic routines eventually evaluate entire analysis space specified by the user. However, the algorithm prioritizes the evaluation across smaller units of the analysis, based on algorithmic and/or user biases, with the goal of converging on exception or patterns faster. Exceptions and patterns may be determined by using combination of statistics and data mining techniques, such as exponential trend, Chi-Squared Deviancies, ANOVA, MANOVA, Cramer's coefficient, entropy, Classical Clustering, SOM, Affinity Analysis etc., depending on the objective of the analysis and the analysis scope definition.

[0117] As the analysis of various crosstabs proceeds, the normal expected behavior is constantly updated and Local, Regional and Global versions of such intelligence are maintained and utilized for generating high exception probability crosstabs and prioritizing the analysis across these crosstabs as has been discussed earlier. In some cases the crosstab definitions along with the prevailing thresholds may be persisted so that it can be reutilized in subsequent analysis runs. An alternative application of the concept of Composite Measures and Binned Ranges is in Ad-

Hoc reporting, which is also utilized for presenting the exceptions and patterns in the If ... Then... form.

[0118] Use of composite measures in the creation of candidate crosstabs helps in overcoming the hard boundaries posed by pre-aggregated structures and offers new data points for analysis. Seed (or Global), regional and local behavioral intelligence of composite measures and user defined biases can be effectively used to guide analysis to converge on anomalies and patterns quickly and yield better insights. Partitioning the analysis scope into smaller units helps with scalability and effective analysis processing. Composite measure definitions and bin ranges thresholds when presented with the exception and patterns helps with ease of understanding.

[0119] During iteration, when an anomalous cellset is detected, a new process (in parallel or for post-processing) can be started to further qualify the anomaly and investigate behavior along previously unconsidered conditions. The process interface will optionally allow for capturing user biases related to prioritization of such new investigative processes and also to define a new scope along previously unconsidered conditions.

[0120] This technique provides a guided path for the statistical techniques to find exceptions and converge on the anomaly fast. However, it is important to note that the ultimate search space is defined by the multi-dimensional database and optionally by user selection. The algorithm evaluates the entire search space eventually; however, process intelligence and/or user biases define priorities of such evaluation. Thus, leveraging the synergy between automated algorithms and the subject matter expert, as represented by the user. The statistical outlook (for the entire population) on a sample data (for performance reasons) is revised intelligently (in local zones) to minimize false exceptions. The results of such analysis could be the identification of local trends within data bins as shown in Figure 18.

[0121] This processing can be done in incremental passes – where new "Normal" relationship is defined based on detected exceptions, which now serve as new sample population.

[0122] For example, crosstab combinations are computed with the goal of converging to potential anomalies first, e.g., along Computers with High Paging and High Memory Availability, so that most prominent Dimensional Members are analyzed First. Such prioritization of binned data is shown in Figure 19, where the data is analysis in the order shown. Thus, the problem definition is divided along multiple parameters with the goal to converge to anomalies fast. While Bin Combinations define scoping based on Crosstab Qualifications, simultaneous use of prominent members (along dimensions/levels etc., e.g., computers with

higher memory availability, further scopes the problem, such that prominent anomalies are detected fast, a result of such an analysis is shown in Figure 20, where the anomaly is shown in the hatched circle.

[0123] Alternately, combinations of statistical tools can be utilized to detect variations/anomalies: for example, exponential trend, Chi-Squared Deviancies, Cramer's coefficient, entropy, Classical Clustering, SOM *etc.*

[0124] Other statistical techniques that can be used with the present methodology including Tendency Analysis, where the main idea is to be able to a collection of "profile" curves representing tendencies in the data. To "detect" significant changes for each event (combination of dimension members), it is enough to detect changes in curve in the tendency plots.

[0125] As a summary, the user can select the dimensions he/she wants to use as an event for the analysis. The algorithm generates several queries to get the tendency curve and a simple technique will be used to detect significant changes. If those changes are significant, the algorithm reports this combination of members to the user. This "Tendency Analysis" can provide information only about when a deviation in the tendency occurs.

[0126] Besides Tendency Analysis, condition analysis can be used. If the user wants to find out which conditions are tending to fluctuate, then the method will need to detect all combinations that move their tendency one way or the other. Thus, condition analysis provides information about where a change in the tendency occurs. The method can ask the user to select if he/she wants a positive or negative tendency and we can show all conditions that have the requested tendency.

[0127] One important aspect of the methods of this invention relates to determining bin sizes, and what approach to take to iterate across Bins, *e.g.*, independent bins (across each composite measure) or process bins one at a time as shown in Figures 21A&B.

[0128] The one bin at a time approach (as opposed to considering two Bins - along two composite Measure simultaneously) may minimize the effect of Bad Bins and provide sufficient data points to identify exceptions. Most importantly, Binning allows establishing a context of crosstab definition, which when subjected to statistical evaluation, would yield results in a context that can be readily and easily understood.

[0129] The method also compares seed intelligence on each cycle and displays a screen as shown in Figure 22, if the seed intelligence is confirmed.

[0130] Any patents or publications mentioned in this specification are indicative of the levels

of those skilled in the art to which the invention pertains. These patents and publications are herein incorporated by reference to the same extent as if each individual publication was indicated to be incorporated by reference specifically and individually.

[0131] One skilled in the art will readily appreciate that the present invention is well adapted to carry out the objects and obtain the ends and advantages mentioned, as well as those inherent therein. It will be apparent to those skilled in the art that various modifications and variations can be made in practicing the present invention without departing from the spirit or scope of the invention. Changes therein and other uses will occur to those skilled in the art which are encompassed within the spirit of the invention as defined by the scope of the claims.

CLAIMS

We claim:

1. 1. A method implemented on a computer to alter a dimensionality of a multi-dimensional
2 database hierarchical structure, iteratively and dynamically, comprising the steps of:

3 providing a multidimensional database having a native schema,

4 selecting a plurality of members and at least one measure from the schema,

5 merging at least one of the plurality of members and the at least one measure to form an
6 imaginary schema,

7 where the imaginary schema enhances, increases and/or makes more efficient data
8 processing of the data in the dataset so that qualified data points are made available to various
9 data mining algorithms and where the imaginary schema alters a dimensionality of the database.

1 2. The method of claim 1, wherein the dimensionality is increased or reduced or the
2 dimensionality of any part of the database is increases or decreased.

1 3. A method for detecting data anomalies, exceptions or meta exceptions and/or identifying
2 patterns in base aggregated data and/or automatically and virtually generating data points based
3 on base aggregated data including the steps of selecting at least one multi-dimensional dataset,
4 preferably in the form of an OLAP cube, and at least one measure associated with the data
5 dimension in the dataset; capturing a scope of analysis and constraints from a user; constructing
6 a virtual database schema from a native database schema of the dataset to reduce or expand the
7 dimensionality of the dataset as a whole or in regions of interest, while maintaining the
8 associated measure or producing a composite measure from more than one measure; selecting
9 a limited number of data values from the entire dataset or the part of interest; creating an initial
10 global rule, "seed intelligence" describing the behavior of the measure with respect to the
11 selected, limited number of data values; determining data regions that would violate the initial
12 global rule; prioritizing the regions; searching the dataset for data that satisfies the initial seed
13 intelligence and that falls within the regions forming regional datasets; and reporting the regional
14 datasets.

1 4. A method for detecting data anomalies, exceptions or meta exceptions and/or identifying

2 patterns in base aggregated data and/or automatically and virtually generating data points based
3 on base aggregated data including the steps of selecting at least one multi-dimensional dataset,
4 preferably in the form of an OLAP cube, and at least one measure associated with the data
5 dimension in the dataset; capturing a scope of analysis and constraints from a user; constructing
6 a virtual database schema from a native database schema of the dataset to reduce or expand the
7 dimensionality of the dataset as a whole or in regions of interest, while maintaining the
8 associated measure or producing a composite measure from more than one measure; selecting
9 a limited number of data values from the entire dataset or the part of interest; creating an initial
10 global rule, "seed intelligence" describing the behavior of the measure with respect to the
11 selected, limited number of data values; determining data regions that would violate the initial
12 global rule; prioritizing the regions; searching the dataset for data that satisfies the initial seed
13 intelligence forming a compliance dataset and that falls within the regions forming regional
14 exception datasets; if the regional exception datasets are not null (empty) or do not contain too
15 few data points to support statistical analysis, creating regional intelligence or local intelligence;
16 determining datapoints within each regional exception dataset that represent exceptions to the
17 local intelligence; and reporting the results.

1 5. A method for detecting data anomalies, exceptions or meta exceptions and/or identifying
2 patterns in base aggregated data and/or automatically and virtually generating data points based
3 on base aggregated data including the steps of selecting at least one multi-dimensional dataset,
4 preferably in the form of an OLAP cube, and at least one measure associated with the data
5 dimension in the dataset; capturing a scope of analysis and constraints from a user; constructing
6 a virtual database schema from a native database schema of the dataset to reduce or expand the
7 dimensionality of the dataset as a whole or in regions of interest, while maintaining the
8 associated measure or producing a composite measure from more than one measure; selecting
9 a limited number of data values from the entire dataset or the part of interest; creating an initial
10 global rule, "seed intelligence" describing the behavior of the measure with respect to the
11 selected, limited number of data values; determining data regions that would violate the initial
12 global rule; prioritizing the regions; searching the dataset for data that satisfies the initial seed
13 intelligence forming a compliance dataset and that falls within the regions forming regional
14 exception datasets; if the regional exception datasets are not null (empty) or do not contain too
15 few data points to support statistical analysis, creating regional intelligence or local intelligence;

16 determining datapoints within each regional exception dataset that represent exceptions to the
17 local intelligence; update the initial seed intelligence with the local intelligences properly
18 weighted to form an updated seed intelligence; comparing the updated seed intelligence; if the
19 updated seed intelligence is significantly different from the initial seed intelligence, replacing
20 the initial seed intelligence with the updated seed intelligence; repeating the previous three steps,
21 until there is no significant change between the seed intelligence from the previous iteration and
22 this iteration; and reporting the results.

1 6. The method of claim 5, further comprising the step of determining whether a termination
2 condition has been met, where failure to met the condition would restart the analysis construction
3 of the scope of analysis step and the method steps would be continued until the condition is met.

1 7. A method for finding global and local intelligences quickly including the steps of:
2 capturing an analysis scope, a breakdown the analysis scope into one or more logical units
3 or combinations thereof, optionally specifying constraints on the analysis scope; establishing a
4 seed intelligence from a sample data population, from user input or a combination of data
5 sampling and user input, identifying data regions that represent exceptions to the seed
6 intelligence, establishing local intelligence in each exception region, if non empty, updating seed
7 intelligence with local intelligences or forming a composite intelligence of an updated seed
8 intelligence and local intelligences; testing to determine if the seed intelligence or composite
9 intelligence from the last cycle is significantly different than the seed intelligence or composite
10 intelligence of this cycle, exiting changes are insignificant or returning to the identifying step
11 if significant changes occurred for iteration until convergence is achieved, where after
12 convergence, the method will have constructed a consistent intelligence, seed or composite, for
13 describing the data behavior and will have identified exceptional regions, local intelligence
14 associated with the regions and exceptions to the local rules..

1 8. The method of claim 7, further comprising testing the intelligences to determine if a
2 termination condition has been met.

1 9. A method for constructing an intelligence models including an overall or global
2 intelligence and local intelligences using the methods of claims 3-8, which generates the

3 intelligences from the analysis of data in multidimensional databases, relational or OLAP, and
4 in the use the intelligence model to predict further data behavior.

1 10. A method for constructing libraries of intelligence models, each model including an
2 overall or global intelligence and local intelligences using the methods of claims 3-8, which
3 generates the intelligences from the analysis of data in multidimensional databases, relational
4 or OLAP.

1 11. A method for using the library of intelligence models to classify data behavior and as a
2 tool for predicting the behavior of classified data and in the use the intelligence models to predict
3 further data behavior, where the models are generated by the method of claims 3-8.

1 12. A computer having stored thereon code sufficient to implement the method of any of the
2 claims 1-11.

1 13. A computer readable medium having stored thereon code sufficient to implement the
2 method of any of the claims 1-11.

1 14. A system for finding global and local data patterns and exceptions to both the global
2 pattern and the local pattern, comprising:

3 an analysis scope capture and definition module,

4 a breakdown module for breaking the analysis scope into logical units or combinations
5 of logical units,

6 a seed intelligence module that determines a seed intelligence (global rule) from a limited
7 data selection from the data to be analyzed;

8 a determine exception candidate region module where regions of data which would
9 violate the seed intelligence are identified, prioritized and analyzed inline with the analysis of
10 the seed intelligence guess,

11 a determine local intelligence and identify local intelligence exceptions and compare the
12 local intelligence to the seed intelligence,

13 a create an updated seed intelligence module, where the updated seed intelligence and
14 test the updated seed intelligence against the current seed intelligence and repeat the analysis

15 until the updated seed intelligence and current seed intelligence differ by only an insignificant
16 amount.

1 15. A computer having stored thereon code sufficient to implement the system of claim 14.

1 16. A computer readable medium having stored thereon code sufficient to implement the
2 system of claim 14.

1 17. An analysis wizard including a sequence of windows designed to define an analysis
2 scope, define meta dimensions for construction of imaginary database schema, and to defined
3 user customizations or constraints of the imaginary database schema.

1 18. A computer having stored thereon code sufficient to implement the wizard of claim 17.

1 19. A computer readable medium having stored thereon code sufficient to implement the
2 wizard of claim 17.

Dim A	Dim B	Dim C	Dim D	Dim Measure	Intersection Value
A1	B1	C1	D1	Measure	10
A2	B2	C2	D2	Measure	5
A3	B3	C3	D3	Measure	20
A4	B4	C4	D4	Measure	100
A5	B5	C5	D5	Measure	50
A6	B6	C6	D6	Measure	5
A7	B7	C7	D7	Measure	200
A8	B8	C8	D8	Measure	400
A9	B9	C9	D9	Measure	100
A10	B10	C10	D10	Measure	25

FIG. 1A

Dim A	Dim B	Dim C	Dim D	Dim Measure	Intersection Value
A1	B1	C1	D1	Measure	10
A2	B2	C2	D2	Measure	5
A3	B3	C3	D3	Measure	20
A4	B4	C4	D4	Measure	100
A5	B5	C5	D5	Measure	50
A6	B6	C6	D6	Measure	5
A7	B7	C7	D7	Measure	200
A8	B8	C8	D8	Measure	400
A9	B9	C9	D9	Measure	100
A10	B10	C10	D10	Measure	25

Dim A	Dim B	Dim C	Dim D	Dim Measure	Intersection Value
A1	B1	C1	D1	Measure	10
A2	B2	C2	D2	Measure	5
A3	B3	C3	D3	Measure	20
A4	B4	C4	D4	Measure	100
A5	B5	C5	D5	Measure	50
A6	B6	C6	D6	Measure	5
A7	B7	C7	D7	Measure	200
A8	B8	C8	D8	Measure	400
A9	B9	C9	D9	Measure	100
A10	B10	C10	D10	Measure	25

FIG. 1B

FIG. 1C

Meta Exception Definition (1)

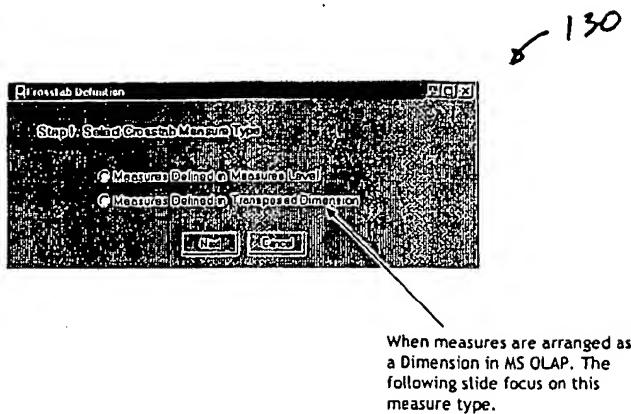


FIG. 2A

Meta Exception Definition (2)

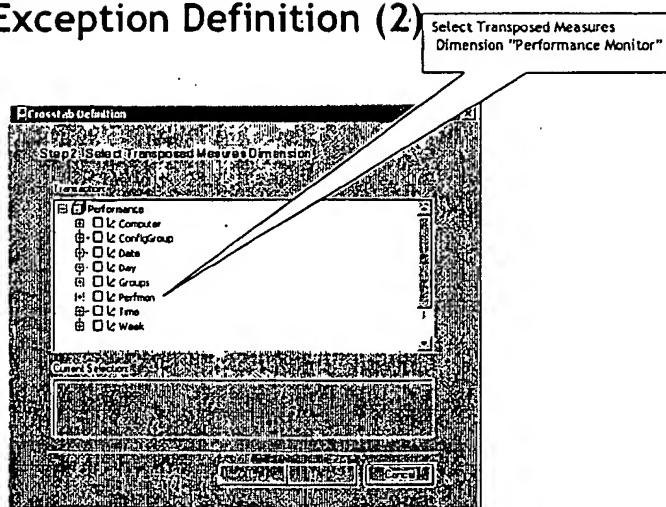


FIG. 2B

Meta Exception Definition (3)

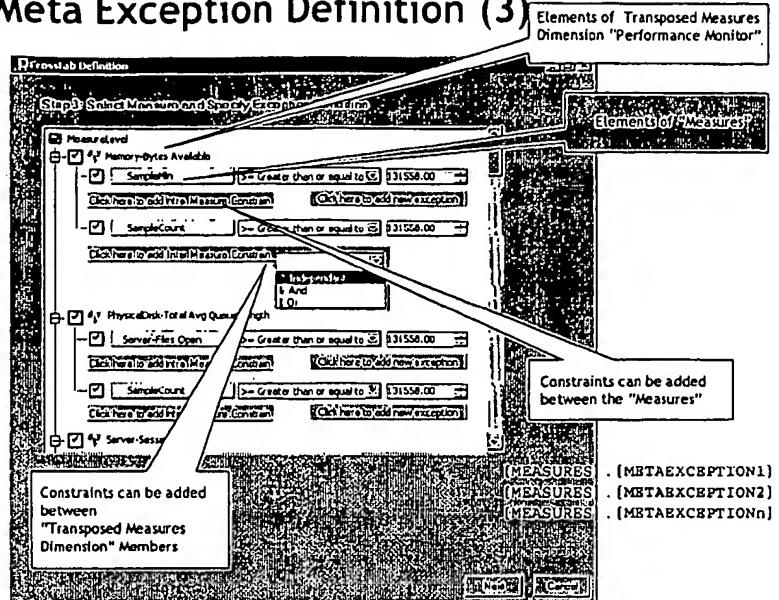


FIG. 2C

Crosstab Definition (4)

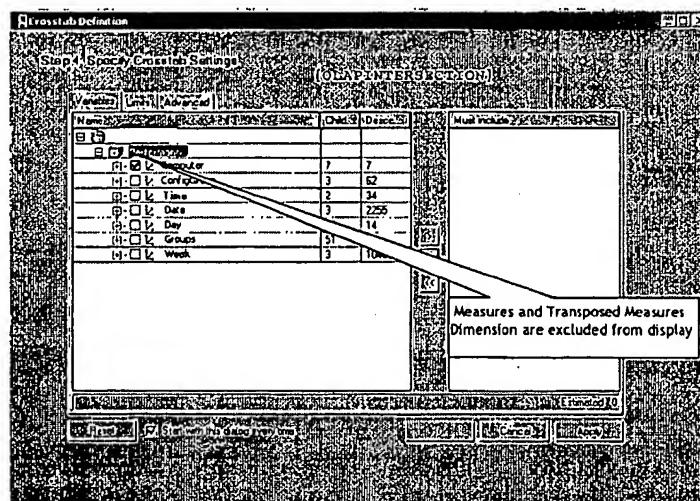


FIG. 2D

Customer	Unit Sales	Customer	Unit Sales
Product	Unit Sales	Customer	Average
Product	Unit Sales	Customer	Unit Sales
Product	Unit Sales	Customer	Average

FIG. 3A

Customer	Education	Customer	Education
Product	Store	Customer	Unit Sales
Product	Store	Customer	Unit Sales
Product	Store	Customer	Unit Sales

FIG. 3B

Customer	Product	Promotion	Promotion
Customer	Product	Product	Product
Customer	Product	502.00	1,548.00
Customer	Product		
Customer	Product	151.00	280.00
Customer	Product		

FIG. 3C

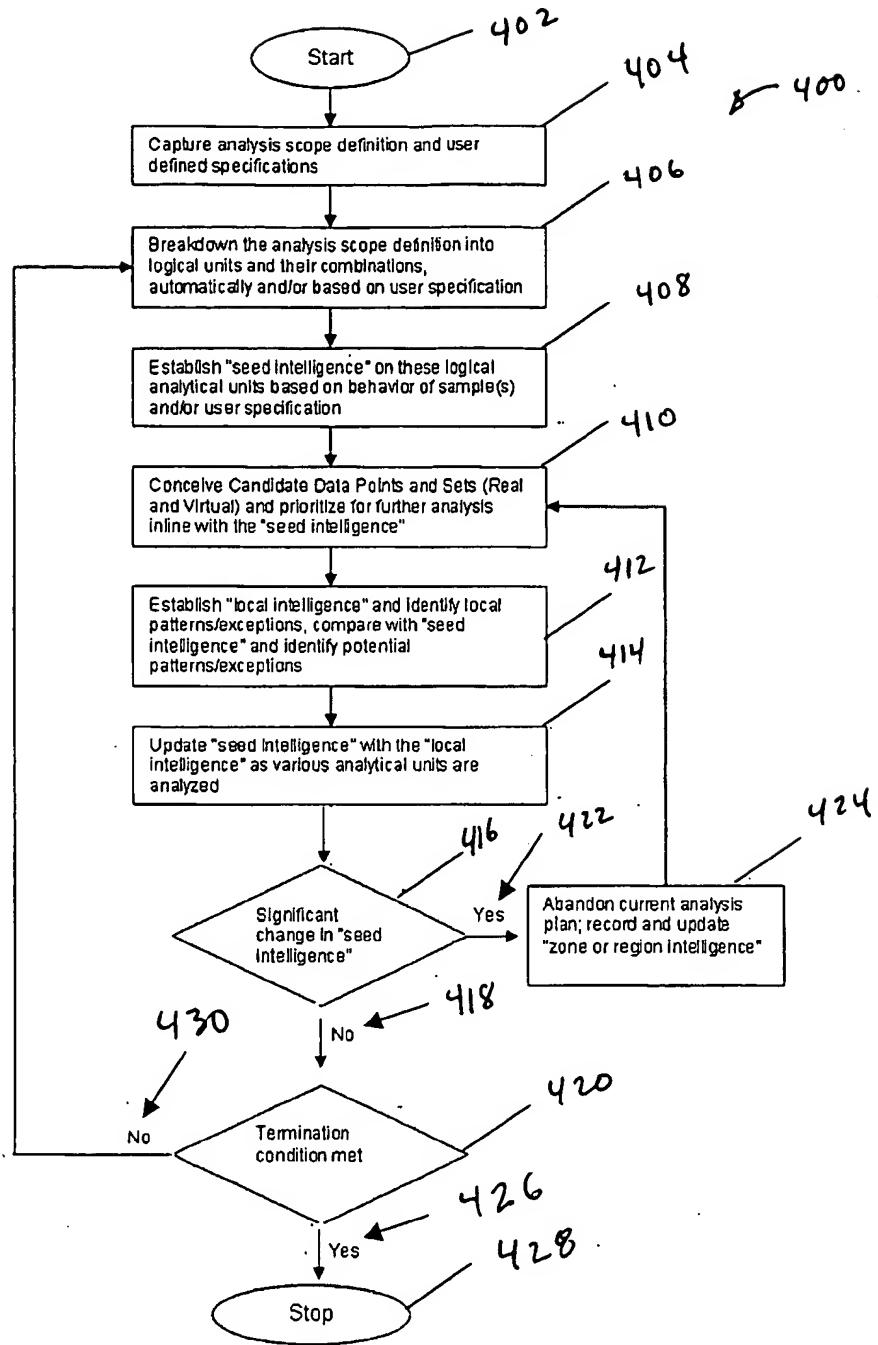


FIG. 4

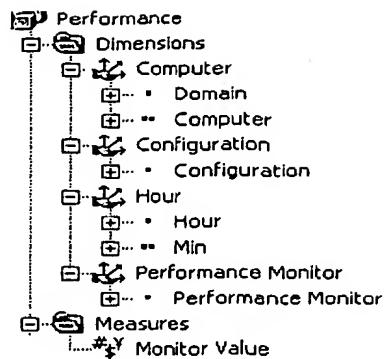


FIG. 5

	JJMemory-Bytes Available	JJMemory-Pages per Sec
JJMonitor Value		JJMonitor Value
JJTOTALS	11,892,278,988,800.00	69,128.77

FIG. 6

PolyVista

Cube

Cube	Children
Performance	5
Computer	2
Configuration	1
Hour	23
Performance Monitor	1
Memory-Bytes Available	0
Memory-Bytes Committed	0
Memory-Reads per Sec	0
Memory-Writes per Sec	0
Memory-Pages per Sec	0
PhysicalDisk:Avg Queue Length	0
PhysicalDisk:Avg Seconds per Read	0
PhysicalDisk:Total Bytes per Sec	0
PhysicalDisk:Total Writes per Sec	0
Process:Total Pagefile Bytes	0
Redirector:Total Bytes per Sec	0
Server:Total Bytes per Sec	0

Run Apply < Back Next > Cancel

FIG. 7

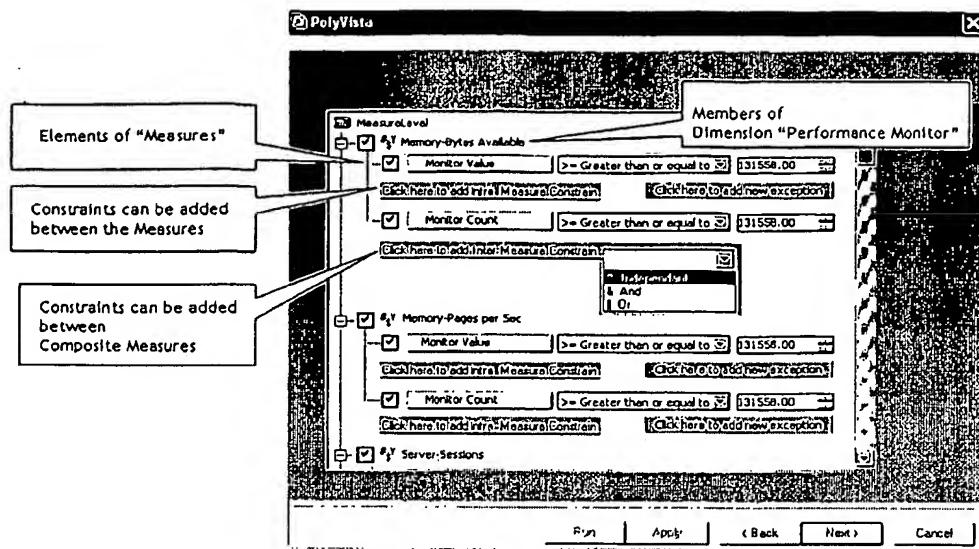


FIG. 8

Memory-Bytes Available_Monitor Value	Memory-Pages per Sec_Monitor Value
11,892,278,988,800.00	69,128.77

F

FIG. 9

FROM

- Performance**
- Measures**
- Computer**
- Configuration**
- Hour**
- Performance Monitor**

To

- Performance**
- Computer**
- Configuration**
- Hour**
- Performance Monitor**

F

FIG. 10

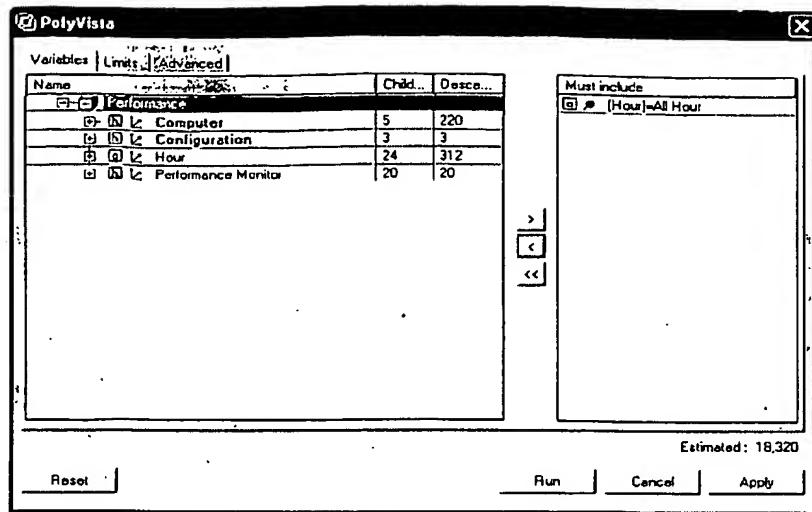


FIG. 11

Correlation_Hours	Correlation_Computer	Correlation_Configuration	Correlation_Overall
-0.92	-0.32	1.00	-0.26

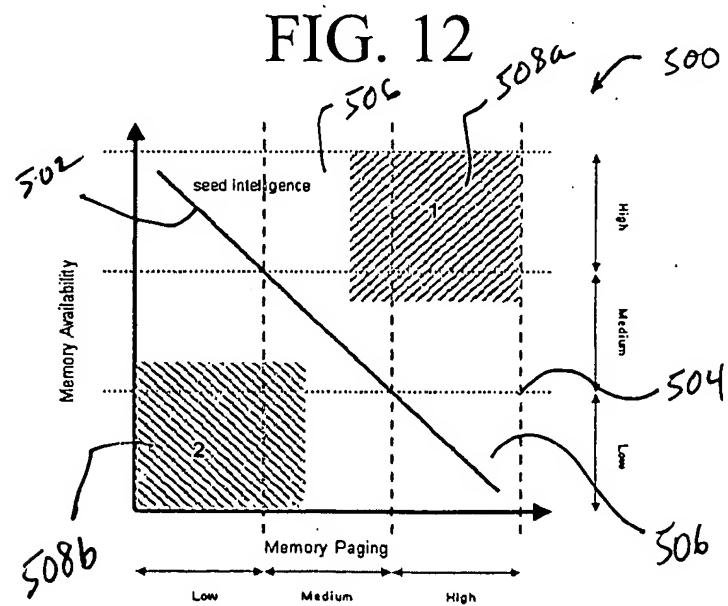


FIG. 13

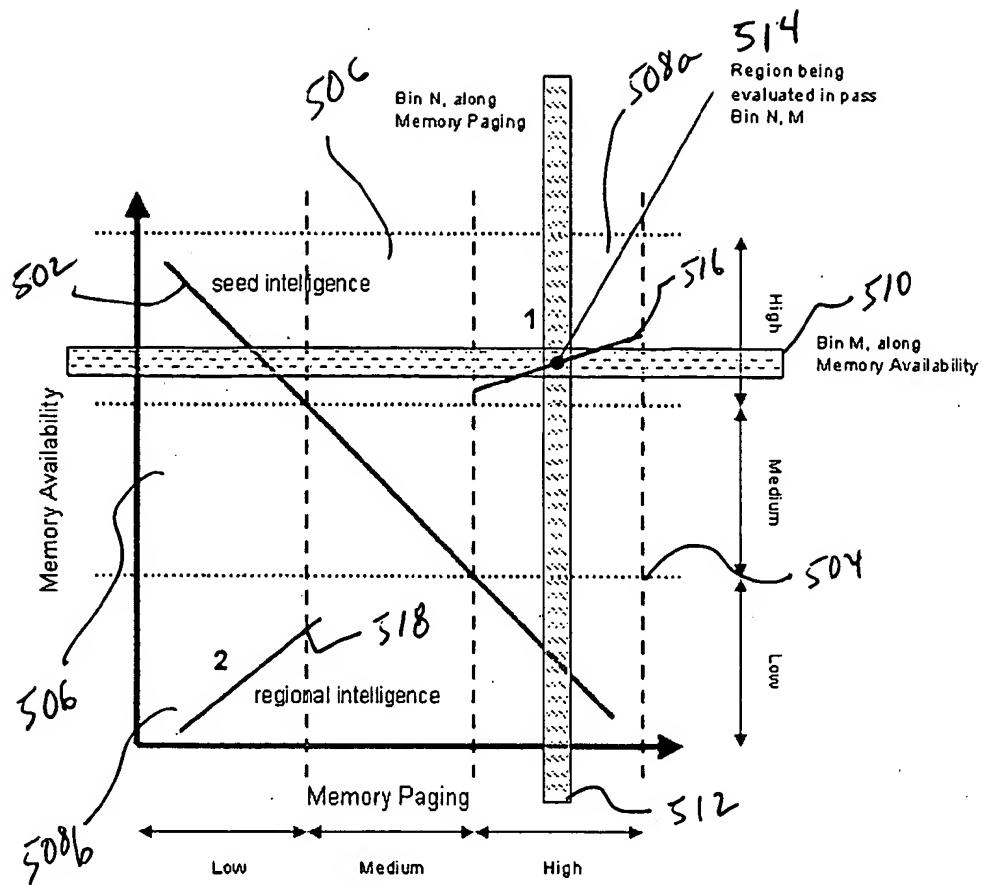


FIG. 14

		Memory-Pages per Sec_(0 To 5)	Memory-Bytes Available_(30 To 60)
Computer176	CG1	18	47.61
		19	45.79
		20	45.79
		22	42.37
		23	58.62
			50.06
Computer198	CG1	17	53.82
		18	56.93
Computer211	CG3	6	58.27
Computer249	CG3	23	55.66
Computer018	CG2	10	56.79
		13	58.63
Computer032	CG2	23	58.66
		0	55.90
		1	57.24

FIG. 15

PolyVista

Iteration: 41 of 100

(Performance Monitor) [Memory-Bytes Available_Monitor Value (0 To 1053206118)], (Performance Monitor) [Memory-Pages per Sec_Monitor Value (764 To 5551)]

Computer	CG1	CG2	CG3	CG4	CG5	CG6	CG7	CG8	CG9	CG10	CG11	CG12	CG13	CG14	CG15	CG16	CG17	CG18	CG19	CG20	CG21	CG22	CG23	CG24	CG25	CG26	CG27	CG28	CG29	CG30	CG31	CG32	CG33	CG34	CG35	CG36	CG37	CG38	CG39	CG40	CG41	CG42	CG43	CG44	CG45	CG46	CG47	CG48	CG49	CG50	CG51	CG52	CG53	CG54	CG55	CG56	CG57	CG58	CG59	CG60	CG61	CG62	CG63	CG64	CG65	CG66	CG67	CG68	CG69	CG70	CG71	CG72	CG73	CG74	CG75	CG76	CG77	CG78	CG79	CG80	CG81	CG82	CG83	CG84	CG85	CG86	CG87	CG88	CG89	CG90	CG91	CG92	CG93	CG94	CG95	CG96	CG97	CG98	CG99	CG100	CG101	CG102	CG103	CG104	CG105	CG106	CG107	CG108	CG109	CG110	CG111	CG112	CG113	CG114	CG115	CG116	CG117	CG118	CG119	CG120	CG121	CG122	CG123	CG124	CG125	CG126	CG127	CG128	CG129	CG130	CG131	CG132	CG133	CG134	CG135	CG136	CG137	CG138	CG139	CG140	CG141	CG142	CG143	CG144	CG145	CG146	CG147	CG148	CG149	CG150	CG151	CG152	CG153	CG154	CG155	CG156	CG157	CG158	CG159	CG160	CG161	CG162	CG163	CG164	CG165	CG166	CG167	CG168	CG169	CG170	CG171	CG172	CG173	CG174	CG175	CG176	CG177	CG178	CG179	CG180	CG181	CG182	CG183	CG184	CG185	CG186	CG187	CG188	CG189	CG190	CG191	CG192	CG193	CG194	CG195	CG196	CG197	CG198	CG199	CG200	CG201	CG202	CG203	CG204	CG205	CG206	CG207	CG208	CG209	CG210	CG211	CG212	CG213	CG214	CG215	CG216	CG217	CG218	CG219	CG220	CG221	CG222	CG223	CG224	CG225	CG226	CG227	CG228	CG229	CG230	CG231	CG232	CG233	CG234	CG235	CG236	CG237	CG238	CG239	CG240	CG241	CG242	CG243	CG244	CG245	CG246	CG247	CG248	CG249	CG250	CG251	CG252	CG253	CG254	CG255	CG256	CG257	CG258	CG259	CG260	CG261	CG262	CG263	CG264	CG265	CG266	CG267	CG268	CG269	CG270	CG271	CG272	CG273	CG274	CG275	CG276	CG277	CG278	CG279	CG280	CG281	CG282	CG283	CG284	CG285	CG286	CG287	CG288	CG289	CG290	CG291	CG292	CG293	CG294	CG295	CG296	CG297	CG298	CG299	CG300	CG301	CG302	CG303	CG304	CG305	CG306	CG307	CG308	CG309	CG310	CG311	CG312	CG313	CG314	CG315	CG316	CG317	CG318	CG319	CG320	CG321	CG322	CG323	CG324	CG325	CG326	CG327	CG328	CG329	CG330	CG331	CG332	CG333	CG334	CG335	CG336	CG337	CG338	CG339	CG340	CG341	CG342	CG343	CG344	CG345	CG346	CG347	CG348	CG349	CG350	CG351	CG352	CG353	CG354	CG355	CG356	CG357	CG358	CG359	CG360	CG361	CG362	CG363	CG364	CG365	CG366	CG367	CG368	CG369	CG370	CG371	CG372	CG373	CG374	CG375	CG376	CG377	CG378	CG379	CG380	CG381	CG382	CG383	CG384	CG385	CG386	CG387	CG388	CG389	CG390	CG391	CG392	CG393	CG394	CG395	CG396	CG397	CG398	CG399	CG400	CG401	CG402	CG403	CG404	CG405	CG406	CG407	CG408	CG409	CG410	CG411	CG412	CG413	CG414	CG415	CG416	CG417	CG418	CG419	CG420	CG421	CG422	CG423	CG424	CG425	CG426	CG427	CG428	CG429	CG430	CG431	CG432	CG433	CG434	CG435	CG436	CG437	CG438	CG439	CG440	CG441	CG442	CG443	CG444	CG445	CG446	CG447	CG448	CG449	CG450	CG451	CG452	CG453	CG454	CG455	CG456	CG457	CG458	CG459	CG460	CG461	CG462	CG463	CG464	CG465	CG466	CG467	CG468	CG469	CG470	CG471	CG472	CG473	CG474	CG475	CG476	CG477	CG478	CG479	CG480	CG481	CG482	CG483	CG484	CG485	CG486	CG487	CG488	CG489	CG490	CG491	CG492	CG493	CG494	CG495	CG496	CG497	CG498	CG499	CG500	CG501	CG502	CG503	CG504	CG505	CG506	CG507	CG508	CG509	CG510	CG511	CG512	CG513	CG514	CG515	CG516	CG517	CG518	CG519	CG520	CG521	CG522	CG523	CG524	CG525	CG526	CG527	CG528	CG529	CG530	CG531	CG532	CG533	CG534	CG535	CG536	CG537	CG538	CG539	CG540	CG541	CG542	CG543	CG544	CG545	CG546	CG547	CG548	CG549	CG550	CG551	CG552	CG553	CG554	CG555	CG556	CG557	CG558	CG559	CG560	CG561	CG562	CG563	CG564	CG565	CG566	CG567	CG568	CG569	CG570	CG571	CG572	CG573	CG574	CG575	CG576	CG577	CG578	CG579	CG580	CG581	CG582	CG583	CG584	CG585	CG586	CG587	CG588	CG589	CG590	CG591	CG592	CG593	CG594	CG595	CG596	CG597	CG598	CG599	CG600	CG601	CG602	CG603	CG604	CG605	CG606	CG607	CG608	CG609	CG610	CG611	CG612	CG613	CG614	CG615	CG616	CG617	CG618	CG619	CG620	CG621	CG622	CG623	CG624	CG625	CG626	CG627	CG628	CG629	CG630	CG631	CG632	CG633	CG634	CG635	CG636	CG637	CG638	CG639	CG640	CG641	CG642	CG643	CG644	CG645	CG646	CG647	CG648	CG649	CG650	CG651	CG652	CG653	CG654	CG655	CG656	CG657	CG658	CG659	CG660	CG661	CG662	CG663	CG664	CG665	CG666	CG667	CG668	CG669	CG670	CG671	CG672	CG673	CG674	CG675	CG676	CG677	CG678	CG679	CG680	CG681	CG682	CG683	CG684	CG685	CG686	CG687	CG688	CG689	CG690	CG691	CG692	CG693	CG694	CG695	CG696	CG697	CG698	CG699	CG700	CG701	CG702	CG703	CG704	CG705	CG706	CG707	CG708	CG709	CG710	CG711	CG712	CG713	CG714	CG715	CG716	CG717	CG718	CG719	CG720	CG721	CG722	CG723	CG724	CG725	CG726	CG727	CG728	CG729	CG730	CG731	CG732	CG733	CG734	CG735	CG736	CG737	CG738	CG739	CG740	CG741	CG742	CG743	CG744	CG745	CG746	CG747	CG748	CG749	CG750	CG751	CG752	CG753	CG754	CG755	CG756	CG757	CG758	CG759	CG760	CG761	CG762	CG763	CG764	CG765	CG766	CG767	CG768	CG769	CG770	CG771	CG772	CG773	CG774	CG775	CG776	CG777	CG778	CG779	CG780	CG781	CG782	CG783	CG784	CG785	CG786	CG787	CG788	CG789	CG790	CG791	CG792	CG793	CG794	CG795	CG796	CG797	CG798	CG799	CG800	CG801	CG802	CG803	CG804	CG805	CG806	CG807	CG808	CG809	CG810	CG811	CG812	CG813	CG814	CG815	CG816	CG817	CG818	CG819	CG820	CG821	CG822	CG823	CG824	CG825	CG826	CG827	CG828	CG829	CG830	CG831	CG832	CG833	CG834	CG835	CG836	CG837	CG838	CG839	CG840	CG841	CG842	CG843	CG844	CG845	CG846	CG847	CG848	CG849	CG850	CG851	CG852	CG853	CG854	CG855	CG856	CG857	CG858	CG859	CG860	CG861	CG862	CG863	CG864	CG865	CG866	CG867	CG868	CG869	CG870	CG871	CG872	CG873	CG874	CG875	CG876	CG877	CG878	CG879	CG880	CG881	CG882	CG883	CG884	CG885	CG886	CG887	CG888	CG889	CG890	CG891	CG892	CG893	CG894	CG895	CG896	CG897	CG898	CG899	CG900	CG901	CG902	CG903	CG904	CG905	CG906	CG907	CG908	CG909	CG910	CG911	CG912	CG913	CG914	CG915	CG916	CG917	CG918	CG919	CG920	CG921	CG922	CG923	CG924	CG925	CG926	CG927	CG928	CG929	CG930	CG931	CG932	CG933	CG934	CG935	CG936	CG937	CG938	CG939	CG940	CG941	CG942	CG943	CG944	CG945	CG946	CG947	CG948	CG949	CG950	CG951	CG952	CG953	CG954	CG955	CG956	CG957	CG958	CG959	CG960	CG961	CG962	CG963	CG964	CG965	CG966	CG967	CG968	CG969	CG970	CG971	CG972	CG973	CG974	CG975	CG976	CG977	CG978	CG979	CG980	CG981	CG982	CG983	CG984	CG985	CG986	CG987	CG988	CG989	CG990	CG991	CG992	CG993	CG994	CG995	CG996	CG997	CG998	CG999	CG1000	CG1001	CG1002	CG1003	CG1004	CG1005	CG1006	CG1007	CG1008	CG1009	CG1010	CG1011	CG1012	CG1013	CG1014	CG1015	CG1016	CG1017	CG1018	CG1019	CG1020	CG1021	CG1022	CG1023	CG1024	CG1025	CG1026	CG1027	CG1028	CG1029	CG1030	CG1031	CG1032	CG1033	CG1034	CG1035	CG1036	CG1037	CG1038	CG1039	CG1040	CG1041	CG1042	CG1043	CG1044	CG1045	CG1046	CG1047	CG1048	CG1049	CG1050	CG1051	CG1052	CG1053	CG1054	CG1055	CG1056	CG1057	CG1058	CG1059	CG1060	CG1061	CG1062	CG1063	CG1064	CG1065	CG1066	CG1067	CG1068	CG1069	CG1070	CG1071	CG1072	CG1073	CG1074	CG1075	CG1076	CG1077	CG1078	CG1079	CG1080	CG1081	CG1082	CG1083	CG1084	CG1085	CG1086	CG1087	CG1088	CG1089	CG1090	CG1091	CG1092	CG1093	CG1094	CG1095	CG1096	CG1097	CG1098	CG1099	CG1100	CG1101	CG1102	CG1103	CG1104	CG1105	CG1106	CG1107	CG1108	CG1109	CG1110	CG1111	CG1112	CG1113	CG1114	CG1115	CG1116	CG1117	CG1118	CG1119	CG1120	CG1121	CG1122	CG1123	CG1124	CG1125	CG1126	CG1127	CG1128	CG1129	CG1130	CG1131	CG1132	CG1133	CG1134	CG1135	CG1136	CG1137	CG1138	CG1139	CG1140	CG1141	CG1142	CG1143	CG1144	CG1145	CG1146	CG1147	CG1148	CG1149	CG1150	CG1151	CG1152	CG1153	CG1154	CG1155	CG1156	CG1157	CG1158	CG1159	CG1160	CG1161	CG1162	CG1163	CG1164	CG1165	CG1166	CG1167	CG1168	CG1169	CG1170	CG1171	CG1172	CG1173	CG1174	CG1175	CG1176	CG1177	CG1178	CG1179	CG1180	CG1181	CG1182	CG1183	CG1184	CG1185	CG1186	CG1187	CG1188	CG1189	CG1190	CG1191	CG1192	CG1193	CG1194	CG1195	CG1196	CG1197	CG1198	CG1199	CG1200	CG1201	CG1202	CG1203	CG1204	CG1205	CG1206	CG1207	CG1208	CG1209	CG1210	CG1211	CG1212	CG1213	CG1214	CG1215	CG1216	CG1217	CG1218	CG1219	CG1220	CG1221	CG1222	CG1223	CG1224	CG1225	CG1226	CG1227	CG1228	CG1229	CG1230</th

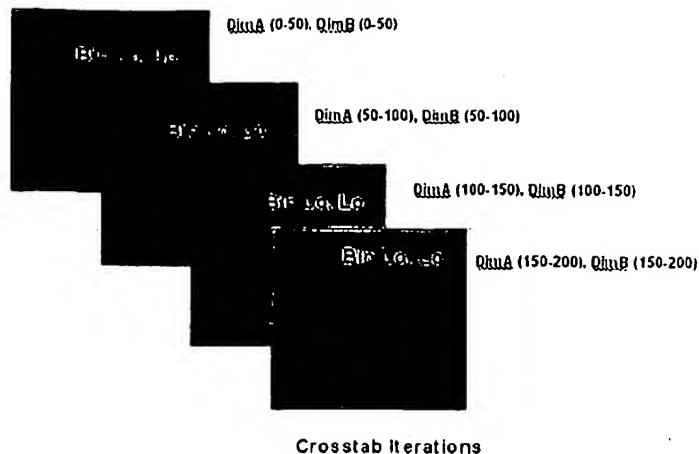


FIG. 19

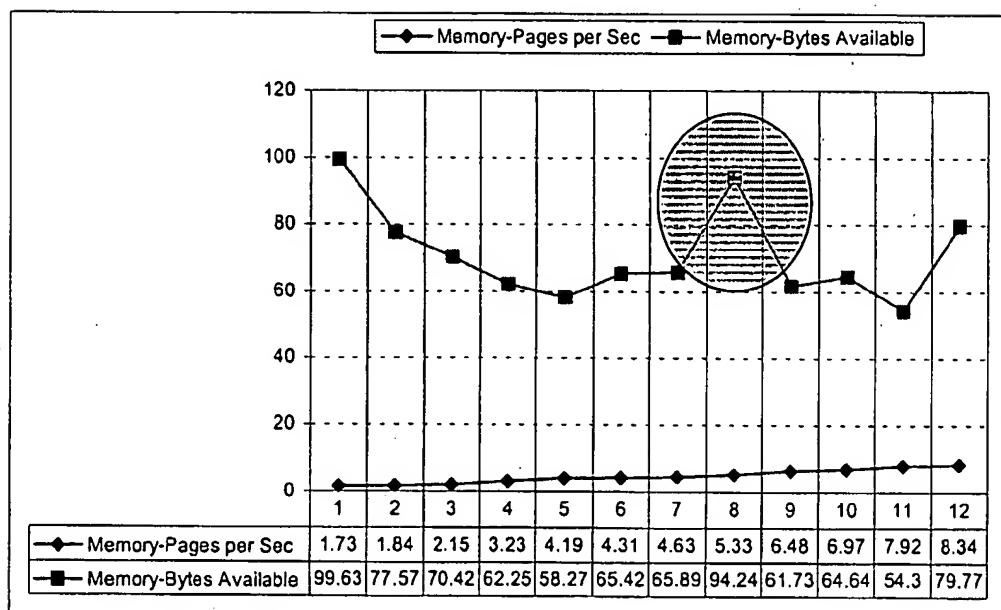


FIG. 20

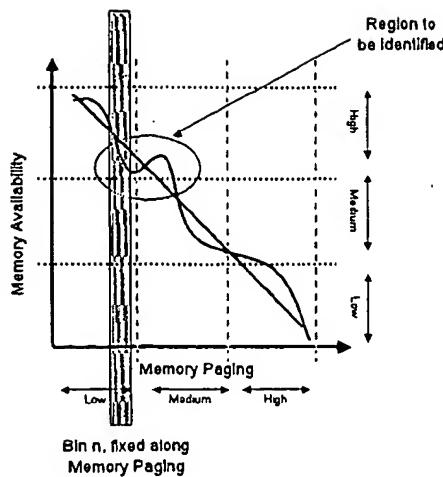


FIG. 21A

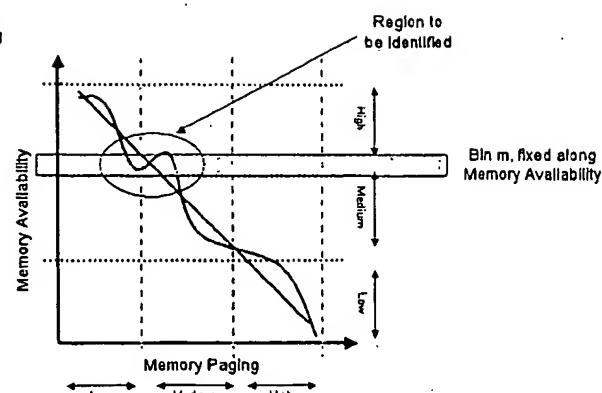


FIG. 21B

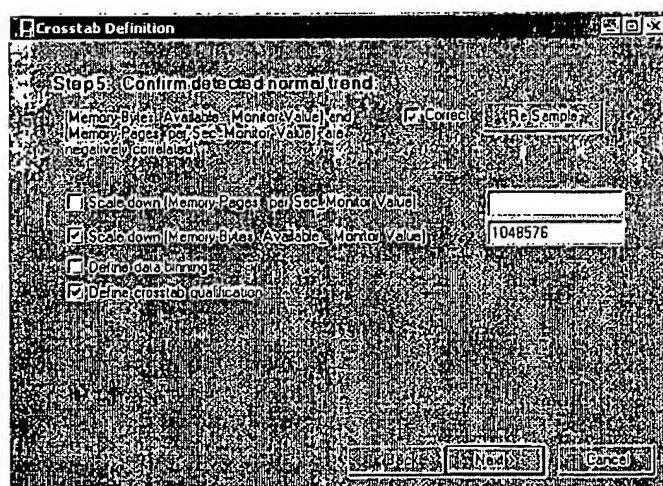


FIG. 22